

LINUX

M A G A Z I N E

THE CHRONICLE OF THE REVOLUTION

GETTING STARTED WITH MySQL

CONFIGURING DNS:

Integrating a BIND
Name Server
with Windows 2000

MANAGED HOSTING:

Evaluating the
Linux-Based Options

ALSO INSIDE:

- More Programming with Threads
- Sharing Data on Dual Boot Systems
- Optimizing Apache



MARCH 2001

\$4.95US \$5.95CAN



03>

0 71486 01105 7

WWW.LINUX-MAG.COM

The American Book Center

1 F 12.90

10% discount for students
teachers and members

TD 2001

More
Linux expertise
per cubicle than
any company
on earth.


IF YOU'RE TIRED OF PAYING FOR OPERATING SYSTEM UPGRADES YOU NEVER ASKED FOR,

YOU HAVE OUR SUPPORT.

With Linux® you can tailor your operating system to the demands of your business —not to someone else's. And with Linuxcare, you now have a single source for everything your organization needs to implement Linux and other open source technologies with confidence. From professional services, to 24X7 support, to training and more, Linuxcare can provide the enterprise-class solutions you're looking for. The Linux revolution is here. And it means business. To learn more, call 888.LIN.GURU (546.4878) or visit www.linuxcare.com.

SUPPORT FOR THE REVOLUTION

LINUXCARESM



No, we don't sit two to the cubicle.

More and more enterprises are turning to Linux and Open Source to gain a competitive advantage. Which often means they turn to us. VA Linux Systems. Why? Because many of the best minds in Linux and Open Source are right here in the cubicles of our Professional Services and Engineering groups. Connecting our customers to the best that Linux and Open Source have to offer. High-powered servers. Network storage. Customized software. Total Internet infrastructure solutions. And our new collaborative development environment for the enterprise. To see what our experts can do for you, call 1-877-VA-LINUX, or visit www.valinux.com.



The World's Linux Leader.

The VA logo is a trademark of VA Linux Systems, Inc. Linux is a registered trademark of Linus Torvalds. ©2001 VA Linux Systems, Inc. All rights reserved.

CONTENTS

MARCH 2001

FEATURES

mySQL

HEY! LEGGO MySQL! 38

Installing, Configuring, and Using MySQL

By Steve Suehring



MySQL is fast becoming the most popular open source database solution available on Linux, and some say it's Linux's next killer app. Once you start to work with it, you'll understand why.

46 OUTSOURCING LINUX

Evaluating Your Linux-Based Managed Hosting Options

By Kevin Railsback

We look at the offerings of eight providers of Linux-based managed hosting services.



SPARKING THE EMBEDDED REVOLUTION 58

Lineo's Bryan Sparks Wants Linux to Go Anywhere

By Robert McMillan

Does Linux really spell the end of proprietary embedded operating systems? Lineo CEO Bryan Sparks certainly thinks so — and he's doing everything he can to make it happen.

64 THE TIES THAT BIND

Using BIND Name Servers with Windows 2000

By Cricket Liu

Configuring a standard Linux-based DNS server running BIND to service Windows 2000 clients can be tricky. Here's what you need to know to make it work.

BIND

LINUX
MAGAZINE
THE CHRONICLE OF THE REVOLUTION

CONTENTS

MARCH 2001

DEPARTMENTS

6 Publisher's Statement

8 Letters

24 Newbies

By *Bill McCarty*

The Truth About Text — Part III

32 In the Trenches

By *Frank Hecker*

Co-opetition in the
Open Source World

72 Guru Guidance

By *Aleen Frisch*

In the Beginning — Part II

80 Compile Time

By *Benjamin Chelf*

Threads and Mutexes — Part II

86 Perl of Wisdom

By *Randal L. Schwartz*

Headlines in the News

92 Tech Support

By *Drew Streib*

Backing Up E-mail,
Sharing Data on Dual-boot
Systems, Optimizing Apache

98 Advertisers' Index

100 Shutdown

REPORT FROM THE FRONT 10

- SourceForge.net: 100,000 Served
- PHP Reaches 5 Million Sites
- AIX and Linux Buddy Up
- 2.4 Linux Kernel Released
- NSA Shares Linux Enhancements
- Turbolinux to Acquire Linuxcare?
- Lineo Shelves IPO Plans
- NT Platform Cracked the Most

REVIEWS 14


- 14 IBM ThinkPad A20p: Linux To Go
- 16 SGI 330 Visual Workstation Packs a Punch
- 18 Save Code with Perforce SCCS
- 20 Learn How to speak *Beginning PHP4*

ZONKER'S PRODUCT PICKS 96

- DISKXtender 2000: Infinite Storage
- Mozilla 0.7: Better Browsing
- Snort: Sniffing Intruders Out
- CodeWeaver's Wine: 2001 Is a Good Year
- BlackAdder: Portable Python
- Gimp 1.2: Graphics Goodness Galore
- Storm Strike 1U Server: Servin' Up a Storm
- Visual SlickEdit: Edit in Style

linux management for a new era of business.

Manage unlimited
systems, and all
flavors of Linux.



All through a
single browser.

CALDERA®

VOLUTION™

Keeping pace with the high-velocity growth of the Internet can make **your head spin**. Introducing Caldera® Volution™, an evolutionary Linux® management solution designed specifically for easy Linux network management. Volution was created by the Linux experts, Caldera, for the new "Internet era" of business. That means you get secure, remote management of multiple Linux systems through a single browser. And you can manage any flavor of Linux on as many systems as you like without having to visit each one, saving you a lot of time and money. Not to mention stress.

Volution lets you manage key elements of your Linux network, including:

- Asset management
- Software distribution
- Systems health monitoring
- Printer configuration management



Find out more about the evolution of network management from Caldera. Visit www.caldera.com/volution3 today for product information and for the FREE white paper: "Linux Management in the New Era."



www.caldera.com

A NEW COMPANY FOR A NEW ERA

UNIFYING UNIX WITH LINUX
FOR BUSINESS™

©2001 Caldera Systems, Inc. All Rights Reserved. OpenLinux, Volution, Caldera Systems and associated logos are trademarks of Caldera Systems, Inc. Linux is a registered trademark of Linus Torvalds. All other brand and product names are trademarks or registered marks of their respective companies.

Better, Faster, Cheaper

Who needs Linux?" That's a question that I have heard a lot of people asking over the past several weeks. It is not being asked in a disparaging or negative way — it's not meant to imply that no one needs Linux.

No, the people I hear asking that question are trying to sell Linux-based products and services, and they are asking because they need to know how best to focus their sales efforts.

It's really an interesting question, because depending on who you ask, you're likely to get very different answers. In fact, the answers to that question tend to highlight the differences between the open source software development model and the traditional closed software development model. They also highlight the differences between the commercial and non-commercial interests of the Linux community.

For example, a good number of the developers that participate in the Linux kernel development effort will likely tell you that *they* need Linux. Linux is a great example of the "scratch your own itch" mentality that fuels so many open source projects. After all, Linux was started because Linus wanted a Unix that he could run on his desktop Intel machine.

Many developers continue to participate simply for the sense of satisfaction that they feel for having contributed to the overall effort. Whether or not the software is actually useful to commercial interests is, at best, of secondary importance to many of these developers.

However, the situation is hardly black and white. It's fair to say that many of the core developers share a more complicated view of their role in the development process. After all, many of them work for commercial Linux companies, so they share the interests of their employer's customers, at least to some extent. They may

need Linux because it fulfills their passions as individuals. However, they also recognize that their customers have specific needs, and those needs may not necessarily be shared by everyone else in the open source community.

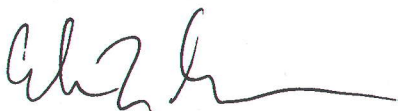
Meanwhile, many companies trying to sell Linux may find that the markets they want to sell to ~~don't~~ feel they particularly *need* Linux at all. This places some vendors of Linux systems and distributions in a particularly complicated position.

All of the commercial Linux players are courting enterprise customers, but enterprise customers often don't particularly care about Linux or open source software *per se*. They simply care about solutions that work and solve their problems better, faster, and cheaper. If those solution happen to be Linux, then that's great. However, if those needs are met by a different platform, they probably could not care less.

Are we as a community as understanding of the needs of our members, both technical and managerial, as we could and should be? Or, do we believe blindly in the principle of enlightened self-interest, fulfilling our own needs and trusting that if we do what's best for us, somehow we'll simultaneously continue to advance the common good?

Rather than further pontificate on the issue, I'd like to put this question to the Linux community, developers and CEOs alike, What do you need from Linux? And what do we as a community need to do in order to deliver on those needs? I'd really like to know.

See you next month,



Adam M. Goodman
President & Publisher

LINUX

MAGAZINE

THE CHRONICLE OF THE REVOLUTION

Editor & Publisher

Adam M. Goodman

Associate Publisher

Lara J. Kisielewska

Managing Editor

Chris Somerville

Reviews Editor

Ross M. Greenberg

Senior Editors

Pete Comas, Matt Welsh

Contributing Editors

Joe "Zonker" Brockmeier, Aleen Frisch,
Bill McCarty, Robert McMillan, Jason Perlow
Kevin Railsback, Paul "Rusty" Russell,
Randal L. Schwartz, Bill von Hagen

Copy Editor

Jason Brainerd

Legal Editor

Steven D. Lavine

Art Director

Rebecca E. Blake

Graphic Designers

Ginny Chu, Mary Norris,
L. Jeanne Veronica

Production Manager

Meredith Rubin

Prepress Technicians

Alberto Lazo, Patty Nohara

Marketing Director

Mike Schaps

National Sales Manager

Eric Derby

Advertising Coordinators

Larissa Jevski, Nancy Prestia

Circulation Director

Melissa A. Witham

Newsstand Distribution

Irwin E. Billman, Ralph J. Perricelli

Marketing Assistants

Greg Aparicio, Andonecia Young

Web Masters

Jon Udell



Accounting Services

Jerome Molejona, Justin Stuhaan

BPA Membership Applied For July 2000

Serve Yourself. SuSE Linux

voted Best Server at
Linux Business Expo —
is available as a fully
integrated platform serving
eMail or Lotus Domino.™



Now available.
Contact solutions@SuSE.com



The Linux Experts

www.SuSE.com 510.628.3380

Prefer full service? Contact SuSE for technical
support and engineering consulting packages:

solutions@SuSE.com



The Secrets of My Success

I recently read "Which Tape Drive Is Best for Your Linux System" (January 2001) and found it very interesting. However, your numbers seem low to me — very low.

For instance, the DLT8000 drive you tested ran at approximately 2 GB per hour. A DLT7000 can easily be run at 20+ GB per hour, so I would expect the DLT8000 to at least match that transfer rate. I have never used Arkeia, so it's possible that those are the best numbers they could sport — but I doubt it.

The following are some techniques I have seen used for trying to get better comparison numbers from backup products, devices, or hardware configurations:

1. Place the files that are to be backed up on a separate drive. This drive should be a different physical unit, not just a logical partition. This can eliminate system contention for disk resources.
2. Back up sets of files (10 KB, 100 KB, 1 MB, 10 MB, etc.). Because of OS overhead, small files are *much* slower than larger files. For example, I've seen 10 KB files backed up at 1.14 MB per second, while 1 MB files ran at 6.26 MB per second.
3. Perform several runs for each file size to ensure that your numbers are consistent.
4. Identify how fast the system can read from disk (e.g., a simple Perl script that measures how fast it can read the data).
5. Identify how fast the system can write to disk.
6. Format the partition prior to creating files for backups.
7. Format the partition prior to

restoring to guarantee that the data will be written contiguously (instead of being fragmented).

8. Remount the partition between test runs when you are doing backups to avoid skewed results from the OS caching files.

9. Perform the benchmarks with easily compressible and difficult to compress files. Do backups with and without compression. A lot of manufacturers choose only to report compressed file backup numbers.

10. If the backup product supports backing up to `/dev/null`, use it to simulate an infinitely fast tape.

I understand your techniques. They may explain your relatively low numbers. However, at least some of the techniques that are listed above will probably yield some different results. Specifically, file size and amount of memory may dramatically improve performance.

Tim Burlowski
VERITAS Software

Why Ignore KDE? GNOME's Not So Great...

In "The Future of Linux" (November 2000), Robert McMillan mentioned GNOME and EAZEL but failed to mention KDE 2. To paraphrase Dan Kusnetzky, VP of System Software Research at IDC, the problem with GNOME is that it is all dressed up with nowhere to go.

Since ditching Abiword, GNOME does not have a word processor. GNOME also aims to completely emulate Windows by including Visual Basic scripts. If I wanted complete conformity of Windows software, why would I bother with Linux? The problem with writing VB code is that it

eats up development time that instead should be spent finishing Abiword and Gnumeric.

The abandonment of Abiword is simply incomprehensible. It's just about complete, and I think it works very well. Why they decided to abandon a word processor that is already integrated into their desktop is a mystery. The decision just doesn't make sense. Also, there is no reason why GNOME could not ship with the Spruce or Balsa MUA in the interim. Instead, the Helix people reinvent the wheel by creating yet another MUA. In the meantime, I can't get e-mail if I use GNOME.

On the other hand, KDE 2 has a fully integrated desktop and is nearing perfection. It may not be as sophisticated as Ximian's GNOME. However, it is complete and works. Software doesn't have to be flashy; it just has to work. End users don't want perfection. Ultimately, they just want predictability.

Jonathan Drews

Correction: In our January 2001 Newbies column, the command we listed for displaying the names of subdirectories was the following:

```
ls | grep '^d'
```


This command will actually list filenames and directories beginning with the letter d. The command should instead be listed as:

```
ls -la | grep '^d'
```


or

```
ll | grep '^d'
```

Linux Magazine welcomes feedback and comments. Letters will be edited for clarity and space; please be sure to include your name and location. Send letters to editors@linux-mag.com.



CONSOLE ACCESS & TERMINAL SERVER



if only
everything
was this
easy.

TS2000

The TS2000 provides secure access to the console ports of servers and networking equipment. Now you can easily manage and troubleshoot your network from anywhere, anytime.

- 32 RS-232 ports on 1U of rack space.
- First Linux-based Terminal Server in the market.
- Packet Filtering, RADIUS, and Secure Shell (SSH).
- Status LEDs on the front panel for all ports.

Also available:
TS1000 (16 ports)



www.cyclades.com

1.888.CYCLADES 1.888.292.5233

510.770.9727

sales@cyclades.com

Fremont, CA



THE LEADER IN
LINUX
CONNECTIVITY



Report from the Front

SourceForge.net: 100,000 Served



They don't have golden arches, but they can still brag about their numbers served. Since SourceForge was launched in January of 2000 more than 100,000 users have registered for what is likely the world's largest Open Source Application Service Provider (ASP). SourceForge experienced a growth rate of 30 percent per month throughout 2000 and now hosts an impressive roster of more than 13,000 Open Source projects including MySQL, Tcl, Python, XFree86, and Squid.

SourceForge provides hosting, mailing lists, CVS and other services for Open Source developers and is a part of the Open Source Developers' Network, owned by VA Linux Systems. <http://www.sourceforge.net>



PHP Reaches Five Million Sites

You could almost say that the growth of PHP has been dynamic. PHP is now being used on more than 5 million domains, according to the PHP Group. PHP is currently growing at a rate of 15 percent per month, faster than the growth rate of the Web itself.




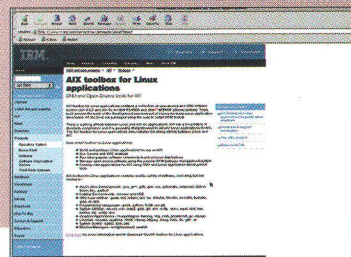
The number comes from the Netcraft survey of Web servers which showed that PHP passed the 5.1 million mark in January.

PHP was developed in 1995 by Rasmus Lerdorf and has become increasingly popular in the last few years. <http://www.php.net>

AIX and Linux Buddy Up

IBM is introducing a toolkit that will let Linux developers use familiar tools to develop programs that run on AIX. The toolkit includes GNOME, KDE, many GNU development tools, popular libraries, and system utilities. Developers will be able to work with C, Guile, Python, and several other languages. The AIX toolkit will also allow developers to utilize package management with RPM.

IBM also proclaimed that both AIX and Linux would be more compatible with the release of AIX 5L. That release will incorporate Linux compatible APIs and header files, which will allow most Linux applications to be compiled as-is and run on AIX. <http://www.ibm.com/servers/aix/products/aixos/linux>



2.4 Linux Kernel Released

The much-anticipated 2.4 Linux kernel was released with little fanfare on January 4th, 2001. Linus Torvalds simply sent a note to the Linux kernel mailing list announcing the kernel's official status. Torvalds acknowledges that the 2.4.0 release probably has a few remaining bugs but that "things don't get better from having the same people test it over and over again. In short, 2.4.0 is out there." The 2.4 kernel series brings improvements in USB support, Symmetric Multi-Processing, a new firewall system, addresses up to 64 GB of memory and many other improvements.

While the 2.4 Linux kernel is officially gearing up, the 2.0 series is winding down. Maintainer Jerry Weinehall announced a "final" update to that series on January 9, making the final version 2.0.39. <http://www.kernel.org>

National Security Agency Shares Linux Enhancements

The National Security Agency (NSA) has chosen Linux as the operating system architecture it will use to build a secure computing environment, and has released a set of security enhancements for the Linux kernel. The enhancements have been released under the GPL and include role-based access control and type enforcement, which determines how users are allowed access to programs or domains.

The enhancements are being made by the Information Assurance Research Office of the NSA in conjunction with NAI Labs, the Secure Computing Corporation and the MITRE corporation. The public is invited to participate in the project, and the NSA is making mailing lists and documentation available along with their security code.

The NSA is responsible for the protection of U.S. information systems, and for procuring foreign intelligence information. <http://www.nsa.gov/selinux>



Turbolinux to Acquire Linuxcare?

Word came from Slashdot first. On January 12, an e-mail reportedly sent to Turbolinux staff by Paul Thomas indicated that a Letter of Intent to merge had been signed by both Turbolinux and Linuxcare. Thomas is the CEO, President, and Director of Turbolinux.

According to C|Net, the talks were confirmed by both companies. Thomas will most likely lead the resulting company, and Turbolinux will keep its board members and add some from Linuxcare.

Both companies have filed for IPOs, though Linuxcare postponed theirs after the departure of CEO Fernand Sarrazat. Both companies have also had to lay off staff, and if the companies merge, there may be further layoffs. <http://www.turbolinux.com>, <http://www.linuxcare.com>

 **turbolinux.**
Powerful Thinking.

LINUXCARE



Lineo Shelves IPO Plans

The stock market that was once so bullish on Linux has recently become a bear. Embedded Linux supplier Lineo filed for an IPO in May of 2000, but is now backing out.

Lineo, which specializes in embedded Linux for set top boxes and other devices, does not plan to layoff any of its more than 300 employees. According to Bryan Sparks, the Chief Executive of Lineo, the company is still growing, but the condition of the Nasdaq led the company to pull the IPO for now.

According to the company's press release, continuing the IPO process would allow competitors access to financial information and business plans as well as "incur unnecessary expense." Rather than continuing to file sensitive information, Lineo requested that the S-1 filing with the SEC be withdrawn.

<http://www.lineo.com>

NT Platform Cracked the Most

Unlike the US Presidential Election, the most cracked server platform race had a clear winner right away — and it was Windows NT. Attrition.org reports that Windows NT was the most compromised Web platform between August 1999 and December 1999 with 56 percent of reported

defacements. While Windows NT led the hall of shame, Linux came in a distant second with all Linux distributions totaling 21 percent of the total. Solaris hosted 8 percent of the cracked sites between August and December. All of the BSD flavors together totaled only 6 percent of the cracked Web sites.

Operating System results were derived from querying the site at the time of the compromise with nmap tools and NetCraft. Red Hat Linux came in first among Linux distributions with just over seven percent of defacements, while "unknown" Linux distributions accounted for a little more than twelve percent of cracked boxes. Statistics were compiled by Attrition.org and are available at the Attrition.org Web site at <http://www.attrition.org/mirror/attrition/os.html>.

Microsoft
**Windows NT
Server**



E-D9228-60000-1

SYM	PRODUCT	DATE	DESIGN GOAL
A	NETSERVER LH 3000	02/23/00	ON [∞]

DUAL INTEL® PENTIUM® III PROCESSORS.
UPGRADEABLE TO INTEL PENTIUM III XEON™ PROCESSORS.

EIGHT HOT-SWAP FANS REGULATE
TEMPERATURE IN SEPARATE
"COOL ZONES"

• INTERNAL DUAL-CHANNEL RAID
CONTROLLER W/OPTIONAL
BATTERY BACKUP

• EASY TO CONVERT FROM
PEDESTAL TO RU RACK

• ALL PWR. SUPPLIES
TO BE HOT-SWAPPABLE.

COLOR-CODED
TOUCH POINTS.

• SERVICE ALL MAJOR
COMPONENTS w/O TOOLS

• ONE-BUTTON DISASTER
RECOVERY (OPT.)

• BACKUP
VIA INTERNAL
SURESTORE T20
OAT 24 X 6I (OPT.)

CAN
WE MAKE
PHYSICAL
CONFIG. EVEN
MORE INTUITIVE?

• IF IT'S CRITICAL,
IT GETS A BACKUP
(FANS, POWER SUPPLIES,
CONTROLLER CARDS, HOST BUS
ADAPTERS, HARD DRIVES, ETC.)

INCLUDE REMOTE
ASSISTANT AND
TOPTOOLS.

• ELIMINATE INTERRUPT PROBS.
(TRY PARALLEL I/O CHANNELS/IMPROVED
INTERRUPT SHARING ALGORITHM)


PRELIMINARY SPECS

- 128 MB SDRAM STD. W/ECC-EXPAND. TO 8 GB
- 256 KB ON-DIE WRITE-BACK CACHE
- UP TO 12 HOT-SWAP ULTRA2 SCSI DRIVES
- EIGHT FULL-LENGTH PCI SLOTS (4 HOT-PLUG)
- INTEGRATED 10/100TX NIC

www.hp.com/go/pc-us



pentium® III



The real genius lies in what you do with it.

Intel, the Intel Inside Logo and Pentium are registered trademarks and Pentium III Xeon is a trademark of Intel Corporation. ©2001 Hewlett-Packard Company. All rights reserved.



IBM Corp. ThinkPad A20p \$3,899
<http://www.ibm.com/thinkpad>

In a Nutshell

Rating:




Pros:

- Huge, beautiful 1400x1050 display
- No problems with Red Hat 7 out of the box
- Affordable compared to other "desktop replacement" notebooks

Cons:

- IBM doesn't mention or support Linux on this model
- Too large and heavy to use on an airplane



Specifications

CPU

- 700 MHz Intel Pentium III processor

RAM

- 384 MB RAM

Hard Drive Space

- 18 GB hard disk

Extras

- 6x DVD-ROM drive
- Three-year warranty

Additional Resources

- http://www.zhlive.ch/zhlive_contents_linux.html
- <http://www.linuxcare.com/labs/certs/pada20p-suse64.epl>
- <http://www.cs.utexas.edu/users/kharker/linux-laptop>

By Alan Zeichick

IBM's high-powered ThinkPad A20p notebook blurs the lines between portable computers and desktop PCs. Based on a 700 MHz Intel Pentium III processor with an 18 GB hard disk, complete with a huge 15-inch 1400x1050-pixel screen, this notebook can hold its own against most desktop systems.

However, at nearly eight pounds, it's too heavy a notebook to carry around and too large to fit on an airplane tray table or balance on your lap. Think of the ThinkPad A20p notebook as a desktop system with the ability to be hauled home, carried to a field office, or lugged to a hotel room whenever it is necessary.

Running Linux

The system reviewed came equipped with a 6x DVD-ROM drive, 384 MB RAM, and Windows 2000 Professional (go figure). IBM's support Web site for this ThinkPad provided no information about upgrading to Linux; Linux isn't even mentioned in a separate "compatibility matrix" document on the IBM site. Thus, we thought Linux wouldn't support all of the notebook's special hardware. We also expected the SVGA+ display panel and Crystal Semiconductor sound chip to give us problems.

Fortunately, none of these technologies provided a challenge for Red Hat Linux 7.0. Our only problem was a crash when Xconfigurator attempted to probe the 16 MB ATI Rage Mobility-M3 graphics chip; however, it installed fine on a second try without the probe, the GNOME desktop appearing automatically in an expanded, somewhat fuzzy 800x600 mode. Surprisingly, re-running Xconfigurator and searching through the screen

types revealed a 1400x1050 display option. When we tried it, the screen lit up in full resolution, which is excellent for Web surfing or having multiple windows open at once.

Our initial concerns about the audio system were also unjustified; the sound coming from the two stereo speakers was tinny but clear, even when playing a CD-Audio disk. Red Hat 7's drivers also detected and correctly installed the DVD player, built-in 10/100 Ethernet card, and V.90 modem. GNOME's battery meter worked with the ThinkPad's battery management and accurately reported the percentage of charge in the Li-Ion battery and predicted how much runtime remained.

Power and Speed

Once up and running, the ThinkPad was a joy to operate. Not only does its screen best most 17-inch monitors in brightness and resolution, but the computer's blinding speed, storage capacity, and full array of ports makes it the equal of all but the latest GHz desktops. In addition, having an on-board Ethernet connector is preferable to having to add a separate PCMCIA- or USB-based NIC. We also found the eraser-shaped "TrackPoint" mouse and smooth keyboard perfect for professional writers.

No Respect

Still, considering how much work the open-source community has done to provide drivers for ThinkPad hardware, it would be nice if IBM acknowledged their effort and provided some technical notes or at least listed Linux on its "compatibility matrix" for this computer. But forget about IBM's obligatory "Built for Windows" stickers — this is one classy notebook that runs Linux with style. **LM**



Impressive

VERITAS NetBackup™ *for Linux.*
Conquer Linux backup.

VERITAS NetBackup *for Linux.* The Value of Leadership.

The #1 provider of backup solutions in America is the natural choice for your Linux environment. With NetBackup *for Linux*, VERITAS offers a true enterprise-level backup and recovery solution. It's advanced, yet intuitive and easy to use; comprehensive, yet remarkably affordable. Then consider the proven performance, scalability and VERITAS quality support, and it's clear that NetBackup *for Linux* delivers value you thought was beyond your reach.

For more information on VERITAS NetBackup *for Linux*, visit www.veritas.com/linuxbackup.


VERITAS™
BUSINESS WITHOUT INTERRUPTION™

1-800-729-7894 code: 8450109AI | www.veritas.com

Copyright © 2001 VERITAS Software Corporation. All Rights Reserved. VERITAS, VERITAS SOFTWARE, the VERITAS logo, Business Without Interruption, VERITAS The Data Availability Company and VERITAS NetBackup are trademarks or registered trademarks of VERITAS Software Corporation in the US and/or other countries. Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

SGI 330 Visual Workstation for Linux
\$6,000–\$6,500
<http://www.sgi.com/workstations/330>

By Jason Perlow

In a Nutshell

Rating:




Pros:

- Unparalleled 3D graphics performance for the price
- No-nonsense driver and graphics library configuration with the supplied SGI Linux Pro Pack
- Supports multiple distros

Cons:

- Pricey
- Limited off-the-shelf Linux OpenGL applications
- Pro Pack only supports older Linux distributions at this time



Specifications

CPU

- Dual 1 GHz processors

RAM

- 512 MB RAM

Hard Drive Space

- 18.2 GB SCSI hard disk

Graphics

- VR7 graphics processor

For many years, SGI (formerly Silicon Graphics Inc.) basically owned the high-end workstation market for systems used to create Hollywood-style Computer Generated Imagery (CGI) and highly complex scientific simulations and modeling. If a movie special effect ever made you ask “How’d they do that?”, chances are good it was done on an SGI machine.

SGI has gone through many trials and tribulations over the last few years, mainly because their lightning-fast, ultra-proprietary MIPS RISC graphics powerhouses were so prohibitively expensive to purchase and support that nobody could afford them, thus affecting SGI’s bottom line.

More Bang for the Buck

After numerous dips in stock price, profit warnings, reorganizations, and executive departures, SGI has reworked their product line to contain more modest computer workstation and server offerings. Their product line is now composed mostly of commodity hardware components and runs on Intel chips and PC-based OSes, such as Windows NT and Linux. Basically, their machines no longer cost both an arm and a leg.

SGI’s latest offering, the 330 Series Visual Workstation, combines their expertise in the high-performance workstation market with more affordable off-the-shelf components, albeit on the high end of “off-the-shelf.” Our test unit boasted twin 1 GHz Pentium III processors, 512 MB of 133 MHz ECC SDRAM, an onboard 16-bit sound chip, onboard Intel Pro 10/100 Ethernet and embedded twin-channel Adaptec 7899 Ultra-160 SCSI, with

a Seagate 18.2 Ultra 160 10,000 RPM SCSI drive.

Pumped Up

To call the 330 a PC on steroids would be a serious oversimplification; this machine has been optimized for producing lightning-fast, super high-resolution 3D graphics, with no compromise. The onboard 64 MB DDR (Double Data Rate) SDRAM 4x AGP graphics card, branded as the VPro VR7, is a highly tweaked version of the nVIDIA high-end Quadro2 Pro 3D graphics processor. It takes full advantage of SGI’s OpenGL 1.2 3D graphics library and is capable of pumping out more than 31 million triangles per second and up to 1 Giga-pixels per second. This is the type of performance you would expect from a machine being used to design vivid space battles in the next Star Wars movie.

The SGI 330 is not just the latest and greatest in buzzword graphics and PC hardware however. To pull it all together, the 330 runs an extremely tweaked Red Hat 6.2, Turbolinux 6.0, or SuSE 6.4 (support for other distributions and newer Linux versions is coming soon). We performed the installation by first installing the base-level Red Hat 6.2 and then the SGI Linux Pro Pack 1.3 for Visual Workstations.

Installation and Setup

The Pro Pack is an entire CD filled with SGI’s fully compliant OpenGL libraries, graphics drivers, XFree86 4.x, a tweaked SMP kernel, and OSS sound drivers. Fortunately, installation was much simpler than on a typical Linux system. There was no hassling with RPMs or installing packages manually. Just mount the CD, and a simple shell script installs every-

thing for you. One reboot later and *voila!* You've got an instant 3D graphics powerhouse. We wish other big-name PC workstation manufacturers, with their so-called Linux-certified systems, would take the care that SGI does by providing an easy way to get all the drivers you need up and running. SGI clearly knows what they're doing when it comes to Linux.

To test the Visual Workstation, we used SPEC's (<http://www.spec.org>) ViewPerf graphics benchmark and the included SGI 3D demos. We were stunned with the fluidity and speed of the animated 3D models at such high resolutions (1600x1200) and deep color depths (24 and 32-bit). Performance figures varied greatly depending on what resolution and color depth we used, but we were very pleased with how well this unit compared with other manufacturers in the \$6,500 (and higher) price range, given existing scores of graphics workstations that are currently disclosed by SPEC.

However, we were unable to see any evidence of performance improvement when we ran the benchmark using both Pentium III processors. This was probably because the benchmark does not completely exploit multiprocessors on Linux 2.2.x systems. We expect this SMP performance issue to clear up considerably when the 2.4.x kernel-based distributions arrive in the second or third quarter of 2001.

Benchmarks aside, unless you roll your own OpenGL applications in a scientific environment or are perverse enough to use a \$6,500 graphics workstation for playing Quake, there are few off-the-shelf OpenGL applications capable of exploiting this hardware. If you want to give it a shot, Blender (<http://www.blender.nl>, reviewed in the June 2000 issue of *Linux Magazine* and available http://www.linux-mag.com/2000-06/pr_blender_01.html) and a couple of others can be found on Freshmeat (<http://freshmeat.net/appindex/X11/Modelling.html>).

Much of this will change, however, when SGI's Alias/Wavefront subsidiary (<http://www.aliaswavefront.com>) releases the Linux version of MAYA 3.0, the entertainment industry 3D graphics workhorse, currently in late beta, in the second quarter of 2001. Avid Technology is currently porting SoftImage's (<http://www.softimage.com>) next-generation 3D graphics animation suite, codenamed "Sumatra," to Linux and has already released version 4.4 of their TOONZ 2D cell animation packages for Linux.


The Last Word

In terms of hardware and software integration, the SGI 330 is a shining example of what every Linux-compliant PC workstation vendor should supply. However, in this case, stellar Linux integration carries a hefty price tag. Still, those of us on the bottom of the food chain, with our roll-your-own white-box Linux PCs, can always dream of graphics nirvana. **LM**

Console Management Over Telnet!

Access Network Serial Console Ports... From Anywhere!

The CMS-16 Console Management Switch provides secure, in-band and/or out-of-band access to RS232 console ports and maintenance ports on UNIX servers, routers, and other network equipment. System administrators can access remote devices in order to change configuration parameters, connect users to restricted ports, collect buffered data, and perform a variety of other control related functions.



UL US CC NEBS-3 LISTED

western telematic inc.



- ✓ Sixteen (16) RS232 DB9 Serial Ports
- ✓ 10Base-T Ethernet Port
- ✓ TCP/IP Security Features
- ✓ Port-Specific Password Protection
- ✓ Non-Connect Port Buffering
- ✓ Modem Auto-Setup Command Strings
- ✓ Co-location Features
- ✓ AC and 48V DC Powered Models

Visit website for complete NetReach™ product line.

(800) 854-7226 • www.wti.com

REVIEWS

Perforce Software's Perforce \$600
<http://www.perforce.com>

In a Nutshell

Rating:

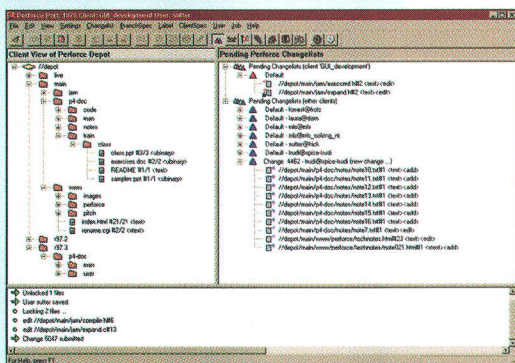


Pros:

- Excellent client GUIs
- Powerful server with many administrative commands
- Source code control is mandatory in complex development environments
- Great customer support

Cons:

- You'll probably need their customer support to get started
- On-line-only documentation irritating
- Mandatory per-client configuration is somewhat complex



Tour de Perforce: With this SCCS you can hack away without fearing someone else will overwrite the work.

Related Resources

Configuration Management Today:

- <http://www.cmtoday.com>

Institute of Configuration Management

- <http://www.icmhq.com>

Newsgroup

- <news://comp.software.config-mgmt>

Open Source Version Control Software

- <http://www.cvshome.org>

Save Code with Perforce

By Bill von Hagen

If you've ever been involved in a programming project that required many developers to work together without the aid of a source code control system, you've probably experienced the joy of having someone overwrite some part of the code you were working on. If this has ever happened to you, you probably now refuse to write code without a source code control package and a loaded pistol by your side. Perforce is an excellent example of the former; for the latter, see the discussion of the `mod_auth_smith_and_wesson` PAM module elsewhere in this issue.

Controlling Code

Source code control systems (SCCS) keep track of changes to a code base, preventing other developers from stepping all over your code. Through "merging," SCCS can also make it relatively easy to blend together multiple changes made to a code base by multiple people without losing anyone's contributions.

Perforce is a client/server source code control system that runs on a large number of platforms. The server typically is where the code store "lives." Clients typically modify only local copies "checked out" from the server. Perforce provides Graphical User Interface (GUI) clients for operating systems such as Windows and Mac OS, and offers an excellent browser-based client called *p4web* for all other platforms (including Linux).

Once you get it working, Perforce is a great piece of software. Unfortunately, that can take a while. The documentation is well-written but hides

some easy getting-started instructions in obscure places (such as within the demonstration program documentation). Luckily, Perforce's customer support is top-notch and patient, no matter how many silly questions you ask (and we asked a lot).

Unfortunately, while the documentation itself was very good, it is only available online, which proved to be extremely irritating. While this may make it possible for software companies to update documentation right up to the last minute, we don't have to like it.

Superior Features

As you might expect with Linux, Perforce faces stiff competition from an excellent piece of free software called CVS (Concurrent Versioning System), which can be compiled on almost any operating system. However, while CVS includes much to recommend it, Perforce definitely earns the money they charge. Perforce's client GUIs are excellent, providing nice hierarchical views of modified files and directories. CVS has free GUIs that make it easy to work with, but they can't touch Perforce's years of dedicated GUI enhancements.

In addition to having fancy GUIs, Perforce has phenomenal technical support. If you're a software developer, you'll be up against a tight deadline at some point, and something will go wrong with your source code control system. It's just the nature of the beast. Perforce may be fairly expensive (especially compared to a product that's free), but you can call them when you have a problem; a patient, pleasant person will help dig you out of the hole you're in.

If your company's future is resting on the successful release of a new product, the kind of service that Perforce provides is priceless. **LM**


Microway Delivers Custom Linux Workstations and Beowulf Clusters for Linux Professionals!

Which Solution is Right for You?

- 833 MHz Dual Alpha
- 500 MHz – 1 GHz Pentium III or Quad Xeon
- 500 MHz – 1 GHz Dual or Single Athlon
- Beowulf Cluster
- Distributed Memory Parallel Processing
- Myrinet Gigabit Ethernet or Dolphin Wulffit
- High Speed, Low Latency Backplanes
- RAID



Microway Screamer™
Dual Alpha UP2000+
833 MHz, 8MB Cache in RuggedRock™
Chassis with Redundant Power Supply



Microway Scalable, 18 Node
36 Processor, Dual Pentium III
Beowulf Cluster with Myrinet and MPI

Satisfied customers have purchased Microway products since 1982. Our cost effective Linux solutions have been chosen by universities, government research labs and ISPs worldwide since 1996. Over the years Microway brand systems have become famous for their quality and price/performance. Simultaneously, our Technical Support staff earned a reputation as one of the best in the industry. This means you can count on our experience to configure a UNIX or Linux workstation or Linux Beowulf cluster solution that is ideally suited to your specific problem. This is why Microway was chosen by Los Alamos National Laboratory to maintain and upgrade their 144-node Alpha Avalon cluster. It is also why both the University of Wisconsin and Rockefeller University chose Microway to build their 100-node dual Pentium III clusters. In addition to being an Intel Product Developer and AMD Athlon OEM, Microway is **Alpha Processor, Inc.'s Top North American Channel Partner.**



"To find a successful supplier of Linux solutions there are three things that you should know: One, the supplier must be proficient in Linux, a master of understanding Linux's superiorities; Two, the supplier must have a clear passion for hardware technology; Three, if you found a supplier that knows both Linux and understands hardware solutions, congratulations you have found Microway! When it comes to Linux and high performance, there's only one way, and that's Microway!"

— Bruce Faust, founder of DigitalScape and Carrera Computers


Microway custom configures **Linux, NT and UNIX workstations, clusters and servers** plus state of the art **RAID** systems. If you need a quality product that is fine tuned and built to last, from a company that will be around to support you for years to come, Microway is **The Number One Choice.**

Find out why over 75% of Microway's sales come from repeat customers. Please call 508-746-7341 for a technical salesperson who speaks your language!

Visit us at microway.com

Microway®

Professional Solutions for Linux Professionals



Research Park Box 79, Kingston, MA 02364
508-746-7341 • info@microway.com

REVIEWS

Beginning PHP4, 800 pgs \$39.99

Publisher: Wrox
<http://www.wrox.com>

Authors: Wankyu Choi, Allan Kent,
Chris Lea, Ganesh Prasad,
Chris Ullman, with Jon Blank and
Sean Cazzell

ISBN: 1-861003-73-0

In a Nutshell

Rating:




Pros:

- Great book for beginners
- Covers basics well
- Practical subject matter

Con:

- Could be better organized



Related Resources

The Home for PHP:

- <http://www.php.net>

Forums, Community, and Links:

- <http://www.phpbuilder.com>

The Development Home for PHP:

- <http://www.zend.com>

Lotssa PHP Code:

- <http://www.phpioneers.com/>

Learn How to Speak PHP

By Joe "Zonker" Brockmeier

Having passed the 5.1 million Web site mark in December, it's pretty clear that PHP is catching on in a big way. If you're a Web developer who is still unfamiliar with PHP, it's time to learn, and *Beginning PHP4* is a good place to start.

The authors of this book do not assume that the reader is a master programmer who just happens to be getting started with PHP. They cover the fundamentals of programming — basic concepts like variables, loops, and arrays — in a way that most beginners will find easy to understand. The authors, however, do expect the reader to know a bit about HTML, but just about anyone interested in PHP will probably have a basic grasp of HTML anyway.

Learning to Program

The authors of this book are very thorough, taking the reader through installation of PHP4 on various Windows platforms, Linux, and other Unices. From there, the authors dive right into demonstrating how to write simple programs with PHP and then progress into some more complicated concepts, such as accessing MySQL databases from PHP code and using PHP to create graphics.

In the "Generating Graphics" chapter, the authors do a particularly good job of breaking down a very complex subject and making it easy to follow. Whether the beginner to PHP is ready to tackle manufacturing graphics on the fly is another question, but they can always skip this chapter. It would have been beneficial to have a chapter on working with databases other than MySQL. PHP works with PostgreSQL, Oracle, mSQL, and several other databases.

Overall, the book is well written and easy to follow. The book's organization is a little confusing, though. Simple concepts like handling e-mail are towards the back of the book after more complicated chapters such as ones on working with MySQL or error handling. The book contains two different chapters on working with clients (browsers), but they are six chapters apart. However, that's really a minor issue since computer books are not usually read front-to-back. The PHP language reference is also helpful, though a bit terse. Since many of the functions are not demonstrated in the book, it would have been useful if the authors explained their usage more completely.

Difficult Concepts Made Easy

Beginning PHP4 really shines when it comes to pulling concepts together. It does so particularly well in the "Case Study" chapter. For example, chapter 17 walks the reader through PHP to create a Yahoo!-like directory that can be dynamically updated and edited. This is a great example tying together earlier material on using PHP functions and working with a MySQL database. It would have been nice to see at least one more case study chapter like this one. We would have particularly liked to see an example that illustrates how to work with the session tracking functions of PHP. In fairness, however, that may be beyond the scope of a beginner's book.

In short, we recommend this book to anyone who wants to get started with PHP4. It's not going to be the only book you'll ever need, but it's a very good starting place. While readers who are already familiar with PHP might want to look for a more advanced book, odds are they'll still find some helpful information in *Beginning PHP4*. **LM**

Dual channel Ultra 160/m SCSI and Dual
NICs for high availability and load balancing

Dual-socket 370 Intel Pentium III
Coppermine processors for speeds
up to 1GHz

3 hot-swappable SCA-2 drives plus
CD-Rom and optional floppy


Optimized 25° Angled DIMM,
supporting high density memory

High-speed chassis blower/cooler,
copper heat sink for optimal cooling

Lancelot Less. Lancelot More!

"Transmeta Corporation is proud to have our Linux/Windows-based computer machines provided by ASL. Their commitment to quality and proven track record is excellent."

— Ray J. Borg, Mgr. IT



ASL
Linux

The Building
Block of your
Internet
Infrastructure



LANCELOT 1200

There is a lot more rackmount here than meets the eye. Our slim form factor and advanced thermal cooling means the Lancelot 1200 takes a whole **lot less** space to do a whole **lot more** work. In fact, the Lancelot 1200 doubles data bandwidth to provide you with high availability for your high-traffic Web-server applications. Lightning fast, the Lancelot 1200 outperforms every other 1U rackmount on the market today. Available with a 4GB high-capacity memory, this is the one rackmount that can handle any size application. Preconfigured with Mandrake or Redhat Linux OS, the Lancelot 1200 comes ready to take on the world for you. Count on the Lancelot 1200 to champion your business: B2B, to ISP, enterprise networking to e-commerce.


POWERFUL. RELIABLE. EXCEPTIONAL.

w w w . a s l a b . c o m
1 - 8 7 7 - A S L - 3 5 3 5



ReDefining Linux Performance

Copyright © 2000 ASL Incorporated. Linux is a registered trademark of Linus Torvalds. Intel, Pentium and Coppermine are property of Intel Corporation. Windows is the property of Microsoft Corporation. All product names are the trade marked, registered trade marked or service marked property of their respective companies.

Sick of Crashing?



Build Your Three-Tier Internet Infrastructure with



The World's Most Reliable Linux Systems


www.penguincomputing.com

©All Rights Reserved. Names, Logos, and Trade Dress appearing on these pages are Registered Trademarks of their respective companies.

Web Server

Colocation facilities are expensive, so your Web servers must have exceptional computing performance to handle large amounts of incoming traffic – yet be slim enough so you can stack multiple units in a rack. With enough nodes to accommodate the current demands of your site, and a slim 1U form factor, Penguin Computing's Relion 120 is the most powerful, flexible and scalable 1U rackmount server in the industry.

Relion 120



- Slim 1U (1.75") form factor
- Up to two Intel® Pentium® III Coppermine processors at up to 933 MHz
- Up to 4 GB of PC133 memory (4 DIMM slots)
- Two integrated Ethernet ports
- Integrated Adaptec Ultra 160 SCSI controller
- Two 3.5" hot-swap 1" SCA LVD drive bays

Application Server

To run your complex business-specific applications like BEA WebLogic Server™, your application servers must have fast processors and large amounts of RAM. To remain in operation, your application servers must also be outfitted with a variety of hardware-based redundant features. Penguin Computing's Relion 220 server is designed specifically for these mission-critical application processes.

Relion 220



- 2U (3.5") form factor
- Up to two Intel® Pentium® III Coppermine processors at up to 933 MHz
- Up to 4 GB of PC133 Memory (4 DIMM slots)
- Two available PCI slots (two 64-bit)
- Six 3.5" hot-swap 1" SCA LVD drive bays
- Up to two hot-swap power supplies
- Up to five Ethernet ports (gigabit optional)

Database Server

Database servers receive large numbers of requests. Programs such as Oracle8i™ that respond to these requests operate best with multiple fast processors, large amounts of RAM and high input / output capabilities. Store and protect your vital business data with Penguin Computing's Magnus 4500.

Magnus 4500



- 4U (7") form factor
- Up to 4 Intel® Pentium® III Xeon™ 700 MHz processors with up to 2 MB of integrated L2 cache
- Up to 4 GB of PC100 memory (16 DIMM slots)
- Six PCI slots: two 64-bit/66 MHz, four 64/32-bit/33 MHz
- Three 300W hot-plug redundant power supplies
- Four 3.5" hot-swap 1" SCA LVA Hard Drives
- Virtually unlimited hard drive expandability using external storage units



NEWBIES

The Truth About Text – Part III

By Bill McCarty

Last month's column introduced you to the venerable Unix editor *vi*. This month's column describes several of *vi*'s more advanced capabilities, which set it apart from other editors. By the end of this column you'll know how to perform editing feats that are quite cumbersome without the help of *vi*.

Editors Within Editors

As explained in our previous *Newbies* column, *vi* is really two types of editors in one. The older *ex* editor has been incorporated into *vi*, and any commands beginning with a colon (:) are passed to *ex* for execution.

What wasn't fully revealed last month, however, is that the *ex* editor, though older than the rest of *vi*, is capable of performing some pretty sophisticated actions. Let's begin by con-

sidering *Table One*, which summarizes some important *ex* commands.

In addition to the commands shown in *Table One*, *ex* handles several other important commands that were examined in last month's column, including the commands to write the current file and exit *vi*.

Deleting Lines

The previous *Newbies* column showed how you can delete the current line by issuing *vi*'s *dd* command. Using *ex* commands, you can delete a line other than the current one or delete a range of lines. To de-

lete a single line, issue the command:

```
:xd
```

where *x* is the number of the line that

Table One: Selected *ex* Commands

Command	Description
d	Delete specified lines
g/pattern/command	Execute specified command on lines containing <i>pattern</i>
g!/pattern/	Execute specified command on lines not containing <i>pattern</i>
m location	Move lines to specified location
r filename	Insert text from specified file
s/pattern/replacement/	Substitute instances of <i>pattern</i> with <i>replacement</i>
t location	Copy lines to specified location
! command	Execute a shell command

The set Command

The *ex* command *set* lets you enable and disable options that control the operation of *vi*. The available options vary somewhat from one version of *vi* to another. *Table Two*, pg. 26, summarizes several of the most common and popular options. To view the currently enabled options, issue the command:

```
:set
```

As an example, you can enable line numbering by issuing the command:

```
:set number
```

Most versions of *vi* let you abbreviate the name of an option. Thus, you can enable line numbering by issuing the following command, which is equivalent to that given previously:

```
:set nu
```


To turn off an option, prefix the name of the option with *no*. For example, to turn off line numbers, issue the command:

```
:set nonumber
```

Some *set* commands require a parameter; you can recognize these by the presence of an equals sign (=) in *Table Two*. What kind of command would require a parameter? Well, let's look at the *TAB* character as an example. The *TAB* advances the cursor by a given number of spaces. You can use *set* to set the number of spaces associated with a *TAB* character to four by issuing the command:

```
:set tabstop=4
```


When NETWORK APPLIANCE Users Switch To RAIDZONE OpenNAS, How Do They Spend Their Savings ... Cars? Yachts? Trips?



RAIDZONE OpenNAS 1TB	\$22,650
Mercedes SLK320	\$40,100
Vancouver 32 Yacht	\$62,000
QE2 104-day World Cruise	\$57,600
	<u>\$182,350</u>

**Network Appliance F760
1 TB
\$205,945**

Say you need a Terabyte of fast, reliable network storage for Windows or NFS. You could buy the NetApp 760 and be quite satisfied. Or, you could buy the RAIDZONE OpenNAS RS2000L, the Mercedes SLK320, the 32-foot sailboat, and the 104-day round-the-world cruise (in a deluxe outside cabin), and still have \$23,000 left over. Which would be more than enough to buy a second TB of reliable RAIDZONE storage for your travel videos.

Come on people, this isn't rocket science. When you choose a RAIDZONE storage solution, you join our customer list alongside Boeing, Lockheed Martin, Motorola, Conexant, Palm Inc., Cisco Systems, The Dana-Farber Cancer Institute, Baylor College of Medicine, Komatsu Mining Systems, and hundreds of others.

RAIDZONE OpenNAS RS2000L File Server \$22,650

- 15 hot-swappable, 80 GB Ultra ATA/100 drives
- RAIDZONE disk array, with hot spare, hot swap, RAID 5 and on-the-fly rebuild
- Dual high performance 933 MHz Pentium III CPUs, 256MB RAM
- Gigabit & 100MBit NICs installed
- Complete browser-based remote administration interface
- Supports NFS, Windows (SMB/CIFS) and Apple clients
- Includes full 1 year warranty, with 3 years on the disks

Every one of these industry leaders has purchased a RAIDZONE OpenNAS storage solution. A few of them have already installed over 10 Terabytes of RAIDZONE storage, with more to come. Not surprisingly, none of them has had any problem deciding what to do with the large amounts of money they've saved.

A full line of RAIDZONE OpenNAS file servers is available. If your storage needs are modest, you can choose the 200 GB, 5-disk desktop cube. At the high end, we have a 45-disk, 3 Terabyte rackmount server with dual Gigabit Ethernet NICs. And every size in between. All designed so that you can get the amount of reliable, redundant storage you'll need tomorrow at a price you can afford today.

For more information, or to order online, please visit our Web site:

www.raidzone.com

or give us a call at:

1-800-388-1896




Figure One: Almost all ex commands have the general syntax shown above.

you want to delete. Be sure to type the colon, which signifies that you're issuing an ex command rather than an ordinary vi command. For example, to delete line 3, issue the following command:

```
:3d
```

At this point, you might be asking, "Great idea, but how am I supposed to know what line number I'm looking at?" It's a very good question. The fact is that it's much easier to use this and other ex commands if line numbering is enabled. If your configuration of vi does not display line numbers by default, see *The set Command* sidebar, pg. 24, to learn how to enable line numbering.

Deleting a range of lines is not much more difficult than deleting a single line. To do so, issue the command:

```
:x,yd
```

where x is the number of the first line you want to delete, and y is the number of the last line you want to delete. For example, to delete lines 1-25, issue the command:

```
:1,25d
```

The symbol \$ stands for the last line of the current file. So rather than typing in the line number, you can delete all lines by issuing the command:

```
:1,$d
```

The symbol % stands for 1, \$, so you

can issue this command even more conveniently as:

```
:%d
```

Copying and Moving Text

The ex commands to copy and move text use a syntax similar to that of the d command. To copy a range of lines to a new location, issue an m command as follows:

```
:x,y mz
```

where the range of lines to be copied is x, y, and the copied lines are to be

inserted following line z. If you want to copy only a single line, you can specify x without the comma and y. For example, if you wanted to copy lines 1-5 and insert them after line 10, you would issue the following command:

```
:1,5m10
```

Similarly, to copy line 5 and insert it after line 10, issue the command:

```
:5m10
```

To move lines rather than copy them, use the t command rather than the m command. You may find it helpful to think of the t as standing for *transfer*. For example, to transfer lines 1-5 to the line following line 10, issue the command:

```
:1,5t10
```

Issuing ex Commands

Most ex commands follow the same pattern as the d, m, and t commands. To see this pattern, refer to *Figure One*. As shown there, the typical ex com-

Table Two: Common vi Options

Option	Description
autoindent	In insert mode, causes vi to automatically indent each line to align with the preceding line.
ignorecase	Causes vi to ignore case when searching.
number	Causes vi to display line numbers.
shiftwidth=width	Defines the width of columns used when autoindent is enabled. Also defines the operation of the shift (< and >) commands.
showmatch	Causes vi to briefly flash the matching (or [when the cursor moves over a) or].
tabstop=spaces	Defines the number of spaces of indentation associated with a TAB character.
wrapscan	Causes vi to search from the beginning of the file when the end of the file is reached.
wrapmargin=width	Causes vi to automatically wrap lines that exceed the specified width.

When it comes to data protection, small and medium sized companies, as well as big companies, need fast, reliable and affordable software.

With over 30,000 users worldwide and more than 150 new companies choosing Arkeia® every month to protect their mission-critical data, Arkeia® is the de-facto standard for network backup with Linux.

Visit our web site at www.arkeia.com and benefit from a free 30 day demo of the full Arkeia® 4.2 Professional Version.

A free version of Arkeia® is also available for home and small business use to protect data on one Linux server and two clients.

Arkeia® ... Data protection for the real world.



"A year ago our goal was simple; we wanted one reliable, inexpensive, centralized backup solution capable of supporting all the different operating systems we used. The time we spent keeping track of all the different software and tapes for our operating systems, and the time it took for some of the software to locate and restore data was growing out of control.

Our first decision was to use a Linux based server because of its robustness, reliability and cost-effectiveness. The next step was to find backup software that would run on Linux and allow us to back up the diverse OS's in our environment. After researching several software solutions, we ultimately chose Arkeia for the following reasons:

- *It is easy to install, configure and use.*
- *The server side runs on Linux and can be controlled locally or remotely via X, Java, or command line interfaces simultaneously.*
- *It supports all major operating systems in our environment.*
- *The client side compression provides minimal impact on our network.*
- *The multiframe capability allows several servers to be backed up simultaneously, thereby decreasing the backup window.*
- *It supported the tape library we had at the time, and it allowed us to upgrade to the library we wanted, which added enterprise-level capabilities including barcode reading.*
- *The way it creates tape segments to store the data on the tapes makes restoring quick and easy.*
- *Each tape contains the source code necessary to restore the tape, no matter the situation.*

We have been using Arkeia for approximately a year and have been very satisfied. Arkeia has made the creation of a centralized backup server a breeze. The issue of having to manage multiple systems with different tapes for the different operating systems is no longer a concern."

Barry D. Trout
Sr. Systems Administrator
Giant Industries, Inc.
Scottsdale, AZ

version 4.2 →

arkeia®

www.arkeia.com

mand has four parts, two of which are optional:

- The colon, which signifies that the command is an *ex* command.
- The range of lines that are the subject of the command. If this part is omitted, the command operates on the current line.
- The command itself.
- The command options, which may be omitted. None of the *ex* commands described so far require any options.

Substituting Text

One of the most powerful *ex* commands is the *s* command, which sub-

stitutes text. To replace the string *hi* with *bye* on the current line, issue the command:

```
:s/hi/bye/
```

If you want to perform this substitution over a range of lines, simply specify the desired range. For example, to perform the substitution over lines 1-100, issue the command:

```
:1,100s/hi/bye/
```

The *s* command supports two especially useful options. By default, the *s* command moves to the next line whenever it replaces text. Thus, a second or subsequent occurrence of the pattern in the same line is ignored. To specify that the command should look for multiple occurrences of the pattern

on each line, use the *g* option. For example, the following command will replace all occurrences of *hi* with *bye* throughout the current file:

```
:%s/hi/bye/g
```

Another useful option is *c*, which causes *vi* to prompt for confirmation before performing a substitution.

The full power of *ex* becomes apparent when you learn that regular expressions can appear in the pattern and replacement text of the *s* command. *Table Three* summarizes some important regular expression characters that we looked at in the first part of this series.

To see what's possible by using regular expressions, consider the following command:

```
:%s/^red/blue/
```

This command replaces the text *red* with *blue*, but does so only when *red* is the first word on the line. As an example of a still more sophisticated command, you might consider the following example:

```
:%s/[0-9][0-9]*/number/g
```

This command searches for lines containing an equals sign followed by one or more digits. When it finds such a line, it replaces the series of digits with the text *number*. This kind of raw power is the reason that so many serious hackers swear by *vi* (or *ex*).

Still More *ex* Commands

While the *s* command is truly powerful (See the *Using Hold Buffers* sidebar, pg. 30), *ex* has other, equally impressive tricks. For example, consider the *g* command (not to be confused with the *g* option we spoke about elsewhere in the text), which is summarized in *Table One*. You can use it to, for example, quickly delete all lines containing the text *comment*:

```
:1,$g/comment/d
```


Table Three: Important Regular Expression Characters

Character	Meaning
.	Matches any single character.
*	Specifies that the preceding expression is optional and need not be matched. Moreover, the expression can be matched indefinitely many times.
\x	Specifies that the character <i>x</i> is understood as an ordinary character, even if it is a metacharacter.
[list]	Specifies a list of characters, any one of which can be used in matching. For example, the expression [0123456789] matches any digit.
[range]	Specifies a range of characters, any one of which can be used in matching. For example, the expression [0-9] matches any digit.
[rangerange]	Used to specify multiple ranges of characters, any one of which can be used in matching. For example, the expression [a-z0-9] matches any lowercase letter or digit.
\<pattern	Matches the specified pattern when it occurs as the beginning of a word
\>pattern	Matches the specified pattern when it occurs as the end of a word
^	Matches the beginning of the line
\$	Matches the end of the line

A GIANT LEAP AHEAD



Cornerstone12g – twelve Gigabit ports, zero packet loss. Only from NPI.



Make sure your data survives Gigabit speed.

With independently verified wire-speed, non-blocking performance, NPI's new Cornerstone 12g Gigabit Ethernet switch prevents today's fastest applications from jamming your network. That's because NPI switches are "Powered by NuWaveArchitecture," your guarantee of 100% data throughput with zero packet loss at incredibly low prices per port. Contact NPI today at 1-800-674-8855 or visit www.npi.com for OEM and VAR opportunities.



The Gigabit Ethernet Company™

1-800-674-8855

NETWORK PERIPHERALS INC.

www.npi.com

Following the `g` with a `!` would reverse the sense of the command, causing it to delete lines *not* containing the text *command*.

Here's another *vi* trick — use the `r` command to insert the contents of a specified file into the current file. For example, issue the command:

```
:r test.txt
```

to insert the contents of the file `test.txt` following the current line.

As a final trick, we present the `!` command, which lets you execute a shell command. For example, to list the files in the current directory, issue the command:

```
:!ls
```

If you want to issue several shell commands, start a subshell by issuing the following command:

```
:!sh
```

You can consult the references given at the end of last month's column to learn about other *vi* and *ex* commands.

Practice Baby, Practice...

Okay, you've now had the cook's tour of *vi*. If you aspire to become a chef, rather than a mere tourist, all that remains is for you to practice using *vi*. If you take the time to do so, you'll likely find that *vi* grows on you, eventually becoming an indispensable part of your computing toolkit. Again, to learn even more about *vi*, remember to consult the sources given in the conclusion of last month's column. Until next month, have fun playing around with these editors.

Bill McCarty is associate professor at Azusa Pacific University's School of Business & Management. He can be reached at bmccarty@apu.edu.

Using Hold Buffers

If you like to cut and paste, you'll be glad to know that the `s` command supports hold buffers, which let you store parts of the matched text and use them in the replacement text. This is an extraordinarily powerful capability. To designate a part of the pattern for which the matched text is to be saved, enclose that part of the pattern between the characters `\(` and `\)`. To insert the stored text into the replacement text, specify the text `\1` for the first or only stored text, `\2` as the second stored text, etc. Up to nine parts of the pattern can be stored and retrieved this way.

As an example, consider the following command that helps you change past tense verbs (ending in *ed*) to present tense verbs (ending in *s*) when they occur at the beginning of a line:

```
:%s/^\(.*\)ed/\1s/
```

Okay, that's a significantly hairy command, no doubt about it. However, if you take it a piece at a time, it's not that difficult to understand. The command works like this:

- The `:` signifies that the command is an *ex* command.
- The `%` signifies that the command is to be applied to every line.
- The `s` signifies that the command performs text substitution.
- The `/` marks the beginning of the search pattern.
- The `^` signifies that, to be matched, the pattern must be found at the beginning of a line.
- The `\(` signifies the beginning of a part of the pattern for which the matching text is to be stored in a hold buffer.
- The `.*` is the text to be matched, which can consist of one or more characters of any type.
- The `\)` signifies the end of a part of the pattern for which the matching text is to be stored.
- The `ed` is part of the pattern and must be matched along with the `.*`.
- The `/` marks the end of the search pattern and the beginning of the replacement text.
- The `\1` signifies that the contents of the first hold buffer are to be inserted into the replacement text.
- The `s` is text with which to replace the matched text.
- The `/` marks the end of the replacement text.

Try your hand at using the hold buffer facility. You'll likely amaze yourself at the editing tasks you can perform by using it.

The Highest Density RackServer Solutions Available!

NAServ Network Attached Storage Filers

NAS 3000

\$1,495

as low as \$69 per month*

- 40GB³ Capacity, Max 328GB³
- 2 x 20GB IDE Hard Drive
- 128MB Memory
- RAID 0/1/5 Striping/Mirroring/Parity
- JBOD Support
- 10/100 RJ45 Ethernet Port
- 1U Rackmount (1.75" H/19" W/23"D)

80GB³ (2 x 40)

120GB³ (2 x 60)

160GB³ (4 x 40)

248GB³ (4 x 62)

328GB³ (4 x 82)

NAS 4000

\$2,795

as low as \$94 per month*

- 160GB³ Capacity, Max 492GB³
- 4 x 40GB IDE Hard Drive
- 256MB Memory
- Hardware RAID 0/1 (4x Faster than Software)
- RAID 0/1/5 Striping/Mirroring/Parity & JBOD
- Optional Hot-Swap Drive Bays (Max 4 Drives)
- 2 x 10/100 RJ45 Ethernet Ports (up to 4)
- 1U Rackmount (1.75" H/19" W/23"D)

\$1,795

\$1,995

\$2,295

\$2,695

\$2,995

248GB³ (4 x 62)

328GB³ (4 x 82)

410GB³ (5 x 82)

492GB³ (6 x 82)

\$3,495

\$3,995

\$4,995

\$5,995

E.Serv 1U RackServers (Single & Dual CPU)

1U

\$895

as low as \$35 per month*

- Model 01633-110
- Intel Celeron 633 MHz CPU
- 128MB Memory, Max 512MB
- 10GB Ultra ATA/66 IDE Hard Drive
- Integrated 10/100 Fast Ethernet NIC
- Integrated SVGA Video
- Red Hat Linux 7.0 w/Browser Manager

Pentium III 700 01700-110 \$995

Pentium III 800 01800-110 \$1,195

Pentium III 866 01866-110 \$1,295

Pentium III 933 01933-110 \$1,595

Upgrade to 256MB Add \$99

Upgrade to 30GB HD Add \$49

V-Raptor 1U Dual

\$1,695

as low as \$76 per month*

- Model 12750-230
- Intel Pentium III 750 MHz CPU (Dual Ready)
- 128MB Memory, Max 1.5GB
- 30GB Ultra ATA/66 IDE HD
- Integrated Intel 10/100 Fast Ethernet NIC
- Optional 2nd & 3rd Intel 10/100 Ethernet
- Red Hat Linux 7.0 w/Browser Manager

1 x Pent III 866 01866-130 \$1,895

2 x Pent III 750 01750-130 \$1,995

2 x Pent III 866 01866-130 \$2,395

2 x Pent III 933 01933-130 \$2,795

2 x Pent III 1GHz 011000-130 \$3,295

Upgrade to 256MB Add \$99

DUAL CPU READY

T-Rex 1U Dual/SCSI

\$2,895

as low as \$105 per month*

- Model 11800-S118
- Intel Pentium III 800 MHz CPU (Dual Ready)
- 256MB ECC Memory, Max 4GB
- 18.2GB Ultra 160 SCSI Hard Drive
- Int Adaptec Ultra 160 SCSI Controller
- 2 x Integrated Intel 10/100 Fast Ethernet NIC
- Red Hat Linux 7.0 w/Browser Manager

1 x Pent III 866 01866-S118 \$2,995

1 x Pent III 933 01933-S118 \$3,295

2 x Pent III 866 01866-S118 \$3,495

2 x Pent III 933 01933-S118 \$3,895

2 x Pent III 1GHz 01100-S118 \$4,395

Upgrade to 256MB ECC Add \$399

DUAL CPU READY

Dual CPU 1U under \$2,000!

Plus 1U RackServers from under \$900!

Starting at \$895 & \$1,695 respectively, the E.Serv 1U & 1U V-Raptor are the lowest priced 1U and Dual CPU Ready 1U server solutions in the industry. Plus, the E.Serv 1Us save precious colo space by fitting up to 42 E.Servs in a 19" rack.

Standard Features include Pentium III Motherboard, Integrated 10/100 Fast Ethernet Adapter, Built-in SVGA Graphics, PS/2 Keyboard & Mouse, USB, Serial, Parallel Ports, 1U Rackmount Aluminum Chassis (2 HD, 1 CD/DVD Drive Bays), 230 Watt 1U Power Supply, Einux Compact Cooling System, Red Hat Linux 7.0 OS w/Einux Browser Server Manager. Also supports Unix, Solaris, Windows NT/2000 OS. Dimensions: 1.75" H x 19" W x 23"D.

- 1U Rack height (1.75" H)
- Fit up to 20TBs³ in a Rack
- From 1.1 cents per Megabyte
- RAID 0/1/5 & JBOD Support
- Easy-to-use NAS Manager sets up in less than 15 minutes
- Easily manage user and group access and security
- Auto detects and configures to your network
- Supports Windows, Unix, Linux & Mac clients
- Remotely manage over Internet with standard web browser
- Optional hot-swappable drives (Model 4000)

E.Serv 2U RackServers (Single & Dual CPU)

2U Dual Ready

\$1,395

as low as \$52 per month*

- Model 21800-120
- Intel Pentium III 800 (Dual Ready)
- 128MB Memory, Max 1.5GB
- 20GB IDE Hard Drive (up to 4 HD)
- Integrated 10/100 Fast Ethernet NIC
- Red Hat Linux 7.0 w/Browser Manager

Pentium III 866 21866-120 \$1,495

Pentium III 933 21933-120 \$1,695

Pentium III 1GHz 211000-120 \$1,995

Upgrade to 256MB Add \$99

Upgrade to 30GB HD Add \$49

DUAL CPU READY

2U Dual

\$1,695

as low as \$59 per month*

- Model 21750D-140
- 2 x Intel Pentium III 750 CPU (Dual)
- 128MB Memory, Max 1.5GB
- 40GB IDE Hard Drive (up to 4 HD)
- Integrated 10/100 Fast Ethernet NIC
- Red Hat Linux 7.0 w/Browser Manager

2 x Pent III 866 21866D-160 \$2,095

2 x Pent III 933 21933D-160 \$2,495

2 x Pent III 1GHz 21160D-160 \$2,895

Upgrade to 256MB Add \$99

Upgrade to 60GB HD Add \$99

DUAL CPU READY

2U Dual/SCSI

\$2,495

as low as \$73 per month*

- Model 22866-S118
- Intel Pentium III 866 (Dual Ready)
- 128MB ECC Memory, Max 4GB
- 18.1GB Ultra 160 SCSI Hard Drive (upto 4)
- Adaptec Ultra 160 SCSI Controller
- Integrated 10/100 Fast Ethernet NIC
- Red Hat Linux 7.0 w/Browser Manager

Pentium III 933 22850-S118 \$2,695

Pentium III 1GHz 221000-S118 \$2,995

Upgrade to 256MB ECC Add \$199

Upg to 36GB SCSI HD Add \$499

DUAL CPU READY

Dual CPU 2U under \$1,700!

Affordable 2U RackServer Solutions!

The E.Serv 2U combines super computing power with its dual Intel CPU motherboard and up to 4GB of memory capacity all fitting in just 2U (3.5") of space.

Standard Features include Dual CPU Ready Pentium III Motherboard, Integrated 10/100 Fast Ethernet Adapter, Built-in SVGA Graphics, PS/2 Keyboard & Mouse, USB, Serial, Parallel Ports, 2U Rackmount Chassis (4 HD, 0D, FD Drive Bays), 300 Watt 2U Power Supply, Einux Compact Cooling System, Red Hat Linux 7.0 OS w/Einux Browser Server Manager. Also supports Unix, Solaris, Windows NT/2000 OS. Dimensions: 3.5" H x 19" W x 23"D.

- 2U Rack height (3.5" H)
- Fit up to 21 Servers in a Rack
- Support for Dual CPU
- Supports up to 1GHz CPU
- Up to 133MHz FSB (Bus)
- Easy-to-use Server Appliance sets up in less than 15 minutes
- Remotely manage over Internet with standard web browser
- Standard with Linux, also supports Unix & Windows NT/2000

E.Array 1/4U & 1/2U Cluster Server Arrays

1/2U Server Array

\$2,195

as low as \$79 per month*

- Model 12633-110
- 2 x Intel Celeron 633 CPU
- 2 x 128MB Memory
- 2 x 10GB IDE Hard Drive
- 2 x Integrated 10/100 Fast Ethernet NIC
- Red Hat Linux OS / Einux Browser Manager

2 x Pent 700 12700-110 \$2,395

2 x Pent 750 12750-110 \$2,595

2 x Pent 800 12800-110 \$2,795

2 x Pent 850 12850-110 \$2,995

Upg to 2 x 256MB Add \$299

Upg to 2 x 20GB HD Add \$149

ClusterManager \$799

1/4U Server Array

\$4,995

as low as \$175 per month*

- Model 14633-106
- 4 x Intel Celeron 633 CPU
- 4 x 128MB Memory
- 4 x 6GB IDE Hard Drive
- 4 x Integrated 10/100 Fast Ethernet NIC
- Red Hat Linux OS / Einux Browser Manager

4 x Pent 650 14650-106 \$5,495

4 x Pent 700 14700-106 \$5,995

4 x Pent 750 14750-106 \$6,495

4 x Pent 800 14800-106 \$6,995

Upg to 4 x 256MB Add \$599

Upg to 4 x 10GB HD Add \$299

ClusterManager \$1499

Clustering and data replication

With Einux's optional Cluster Manager™, Einux Servers can be clustered for load balancing allowing for the highest availability without the need for separate, costly load balancing equipment. Einux Servers can be effectively deployed in server farm clusters, which can be easily managed from a remote location using Einux's Cluster Manager™ client where users can easily see server health and availability and manage all services. Best of all, Einux's Cluster Manager™ is OS agnostic and can be implemented on different operating systems (Linux, Unix, Solaris, Windows) in the same network. Also, with Einux's optional Data Mirror Manager™, data can be replicated inside the cluster so users only have to load one set of data, like a website, and the Data Mirror Manager™ replicates the data across the cluster to the other servers. With Einux's Cluster Manager™ and Data Mirror Manager™, Einux's Servers offer the highest availability and redundancy.

Fit up to 168 Servers per Rack!

World's First 4-in-1U and 2-in-1U Server Arrays

Introducing the E.Array, the world's first server arrays based on Einux's patent-pending High Density Cluster Server Array Architecture. Up to 168 individual servers can be installed in a 19" rack making the E.Array perfect for IDCs, ISPs and e-businesses where co-location real estate is precious.

Standard Features include 4 x Pentium III Motherboard, 4 x Integrated 10/100 Fast Ethernet Adapter, Built-in SVGA Graphics, PS/2 Keyboard & Mouse, Serial Ports, Einux Proprietary 4-in-1U Rackmount Aluminum Chassis, 1U Power Supply, Einux Compact Cooling System, Red Hat Linux 7.0 OS w/Einux Browser Server Manager. Also supports Unix, Solaris, Windows NT/2000 OS. Dimensions: 1.75" H x 19" W x 27.5" D.

- 1U Rack height (1.75" H)
- Fit up to 168 Servers (1/4U) or 84 Servers (1/2U) in a Rack
- Supports up to 850MHz CPU
- Easy-to-use Server Appliance sets up in less than 15 minutes
- Remotely manage over Internet with standard web browser
- Standard with Linux, also supports Unix & Windows NT/2000
- Optional ClusterManager for easy high-availability clustering



Sets up in 15 minutes or less!

With Einux's Server Appliance Browser Manager™ System for Linux, you can setup an Einux Server (LX Linux-based models) in less than 15 minutes using any standard Web browser and a TCP/IP Internet or Network connection. The Browser Manager™ allows users to access and maintain Einux servers remotely from anywhere in the world. Best of all, the Browser Manager lowers the Total Cost of Management (TCM) because it makes it so

easy, almost anyone can set up an Einux Server, thus requiring little or no IT experience. Einux's Browser Manager™ comes standard on all Linux-based Einux rackmount servers and arrays. Test drive it yourself at <http://www.einux.com/demo>.

Custom configure your own server at www.einux.com

servicing the next generation

CALL 866.88.EINUX (883.4689)

email: info@einux.com

Toll-Free: 1.866.883.4689 Tel: 626.968.2686 Fax: 626.968.6565

VARs/Resellers: Visit www.einux.com/reseller

*Certain models only. Specs may vary. See specific model for details. *Leasing figures are based on 36 months and are based on approved credit and may vary. *Storage capacities represent raw capacity and may vary according to configuration and partitioning. Not responsible for typographical errors. Prices, configurations and availability subject to change without notice. Call for details. All trademarks are acknowledged. Copyright 2001 Einux, Inc.

IN THE TRENCHES

Co-opetition in the Open Source World

By Frank Hecker

Can there be a happy marriage between for-profit Information Technology companies and volunteer developers? Consider the following situations that occurred over the past few years and the questions they raised:

- Netscape released the source code for the Communicator browser and initiated a new style of open source project — one sponsored and supported by a for-profit corporation. Was this truly an open source project or just a way for Netscape to take advantage of volunteer labor?
- IBM decided to replace its internally-developed Web server technology with the Apache Web server, basing its future Web application products on Apache and seeking to have its corporate developers join the Apache project. Does IBM want to work with the Apache project or exercise power over it?
- Many key developers in the GNOME project quit working as volunteers and joined (or even founded) for-profit companies creating products and services related to GNOME. The project then issued joint press releases with Sun and other major IT vendors. Were GNOME developers still working for a non-corporate project or had GNOME essentially become a vendor consortium that promoted the interests of its corporate participants?

These are just three examples that illustrate the perceived conflicts

of interest that can result when the two worlds of for-profit IT vendors and volunteer open source developers increasingly intermingle.

In some people's eyes this is a simple "community vs. corporation" battle — how the "community" balances spreading the ideals of open source and free software while avoiding a "sellout" to the corporations it is trying to evangelize. In reality, this is not a battle situation at all, but instead a classic example of "co-opetition," where companies that compete for business cooperate on mutually beneficial projects and activities.

"Co-opetition" Grows Up

This concept has become more complex and more interesting in the realm of open source software development, which can include individual developers and non-profit players like the Free Software Foundation (FSF). This can lead to the alliances among

commercial firms and non-commercial development communities. To further cloud the issue, these development communities might include, or even be led by, employees of the commercial firms.


Thus, in the GNOME project, traditional IT vendors (Sun, HP, and IBM), newer "open source vendors" (Red Hat, Eazel, and Helix Code), individual GNOME hackers, and the FSF are working together in an effort to more quickly extend Linux to the desktop.

This new world of "co-opetition" raises a few important questions:

- How can companies sponsoring open source projects work with individual volunteer developers and not leave them feeling exploited?

Netscape pursued two parallel strategies. First, it tried to deliberately distance the open source project from the corporate development effort (e.g., by giving it a separate name, Mozilla, and a dedicated team of supporters, mozilla.org); in this way, participants could feel they were working for the project and not specifically for Netscape.

At the same time, Netscape gave Mozilla contributors access to information and other aspects of the project previously reserved only for employees. Netscape recognized that contributors and employees would not be satisfied without having some measure of control over variables that affected their work. While there is still some tension between the Mozilla project and Netscape's commercial product development efforts, many open source developers have chosen to con-



KEN WESTPHAL



***You don't follow
the crowd***
YOU LEAD THEM.

You're a trend-setter.

You're internet savvy. You make your own rules. You're not only cracking the code, you're writing it. So how is a Web hosting company ever going to keep up with you? ValueWeb™ can by providing you with open source flexibility, operating system stability and enhanced security.

All our hosting packages are backed by our multi-million dollar Data Center, 24/7 technical experts and our unconditional Money-Back Guarantee. With dependable ValueWeb Dedicated Server plans, you'll get the equipment you need, plus round-the-clock server and network monitoring, maintenance and management. Our Cobalt RaQ 3 and VA Linux 2240 2U specials are just two of the many dedicated hosting packages available.

Stay ahead of the competition.

Discover key competitive strategies in the FREE special report from ValueWeb. Get your "On-line Business Success Report" at www.valueweb.com/report

STANDARD HOSTING

\$1995 A MONTH

- 150 MB of disk space
- 20 POP3 e-mailboxes
- 10 GB of data transfer per month
- **FREE** IBM NetObjects™ TopPage Web design software (save \$59)
- Active Server Pages
- **FREE** PowerMailbox™ Starter
- Microsoft® FrontPage® 2000 server extensions
- **Plus** many more features

MERCHANT HOSTING

\$7495 A MONTH

- All of the features of Standard Hosting plus...
- 400 MB of disk space
- 100 POP3 e-mailboxes
- 40 GB of data transfer per month
- **FREE** Equifax™ Secure SSL certificate (save \$345)
- **FREE** Miva™ Merchant™ shopping cart (save \$495)
- **FREE** Credit card merchant setup (save \$299)
- Preprogrammed merchant ID (MID) and terminal ID (TID)
- **FREE** PowerMailbox™ Base (save \$120)
- **FREE** 2 RealMedia streams (save \$96)
- MySQL database with ODBC

DEDICATED SERVER

COBALT RaQ 3

\$195 A MONTH

- 10.2 GB disk space
- 64 MB RAM
- 50 GB of data transfer per month
- Linux operating system
- Microsoft® FrontPage® 2000 server extensions
- File Transfer Protocol (FTP) server
- Apache Web server
- Sendmail
- Secure Socket Layer (SSL)
- Comprehensive site usage reporting

DEDICATED SERVER

VA LINUX 2240 2U

\$1195 A MONTH

- (4) 18 GB 10K RPM Hard Drives in RAID 5
- 1 GB of ECC SDRAM
- 50 GB of data transfer per month
- VA optimized Red Hat™ Linux operating system
- Sendmail
- File Transfer Protocol (FTP) server
- Apache Web server

GET YOUR
**FIRST MONTH
Free**
when you sign up before March 31, 2001

Call 1-800-WE-HOST-U
(1.800.934.6788)

Or sign up on-line at www.valueweb.com/magazine
Please mention Referral Code 1624 to receive this special offer.

Network Solutions®
Alliance
Partner

Intel®
Partner

Microsoft Certified
Solution Provider

Miva
Business Partner

VA LINUX
SOLUTION

COBALT
RAQ 3

e-business

ValueWeb™
Powerful tools make powerful people.™

ValueWeb is a trademark of FloridaNet, Inc. © 2000 FloridaNet, Inc. All rights reserved. All other product names are trademarks of their respective owners. Technical specifications are subject to change. Setup fees may apply.

tribute to Mozilla development and perceive themselves as having a stake in the project's success.

- How can a corporation partner with, and have some influence over, an existing open source project?

IBM's first, and perhaps most important, step was to understand that the Apache project should be thought of as a potential business partner — one to be wooed and negotiated with, not simply a collection of developers to which IBM would add its own people. The result was a set of terms and conditions under which IBM could join the Apache project, with individual IBM developers being “invited in” as they proved their competence and dedication to the project's goals. By forgoing an attempt at short-term dominance within the Apache project, IBM was able to acquire a substantial measure of long-term influence.

rate partners” to simply providing advice, the GNOME community helped ensure that corporate influence over GNOME would be mediated through individuals deemed worthy.

- What will happen next in the ongoing dance of co-opetition among open source projects and commercial companies?

One emerging trend is particularly intriguing; despite business' perceived role as the driving force in the “new economy,” corporations and individuals involved in open source development seem to be rediscovering some traditional functions of politics and government: managing conflict among multiple parties of different interests and beliefs, serving as a counterweight to raw economic power, and increasing the perceived legitimacy of those with authority to make decisions that are binding to everyone.

the power to vote within these projects is not necessarily open to everyone; typically, only those who can claim a significant investment in and contribution to the project carry power. In any case, the heterogeneous nature of the community means that different projects will be free to experiment with different models of organization, just as different states and countries have different governments and constitutions; some projects will have elaborate, formal models of authority, and others will have simple, informal ones. Each project will adapt to its own local conditions and requirements.

Co-opetition: The Next Step in an Ongoing Evolution

This is simply another step in the ongoing evolution of the open source community. Although people worry about the dangers of the open source community being “corporatized,” the open source movement has been corporatized since Netscape's original announcement and business' subsequent enthusiasm for Linux, GNU, and other software. People also worry about the danger of open source being politicized, but political differences have existed since the days of the “GNU Manifesto.” Politics, in one form or another, will continue to play a role in the open source community as long as there are competing interests and disparities in relative power within projects.

We can't go back to some remembered time of innocence; we need to instead move forward and find ways in which individuals and organizations (both for-profit and non-profit) can work together for mutual benefit. Different parties will not always agree and will certainly not always get what they want. However, collaboration is the key to the long-term goal of the open source model prevailing over the proprietary software model, truly changing the way software is developed.

Frank Hecker is Leading Licensing Expert at CollabNet and a mozilla.org member. He can be reached at frank@collab.net.

We can't go back to some remembered time of innocence; we need to find ways for individuals and organizations to work together.

- What issues arise when an all-volunteer project goes “pro”?


As more and more key GNOME developers aligned themselves with commercial GNOME-related ventures (Helix Code, etc.), the GNOME project might have become a battleground for competing corporate interests. This shift might have threatened GNOME's status as an independent project, in perception if not reality.

By initiating formation of a formal GNOME Foundation last August, the GNOME hackers helped ensure that the project would continue as an effort independent of any company or group of companies. Project leaders strengthened the legitimacy of their position by submitting to democratic elections open to all GNOME contributors. Finally, by limiting the formal power of GNOME Foundation “corpo-

There are tradeoffs as well as checks and balances in this new order. By adopting open source licensing, corporations give up some control over distribution and use of their software. They trade this for the benefits of wider adoption and free enhancements of their software. Similarly, by encouraging the formation of separate “.org” bodies for the open source software they release, and by encouraging their developers to work within the rules set by such organizations, corporations are ceding some power over their own development projects. Simultaneously, individual developers may find that working within a formal structure provides their actions with legitimacy as representatives of their employers.

Some may object that democracy has no place in software development projects, stating they should instead be run as pure meritocracies. However,

On-site Linux support.



No matter where you are.

That's right. Linuxgruven delivers support on site, anytime, anywhere—for every application, for any make of hardware. We're there for you 24 hours a day, 7 days a week. No one else can make that promise because no one else has made that commitment.

Linuxgruven fills a long-standing void in the open-source community by delivering vendor neutral Linux support services at the hardware/network, operating system and application levels. This means you can join the OS Revolution and retain the benefits of an established support system.


We can also assist you in:

- Designing and Building your eBusiness Infrastructure
- Porting your existing applications to run on Linux
- Training and Certifying your Technical Staff
- Providing ongoing Administration and Maintenance

Linuxgruven—the only provider of “Anytime, Anywhere” Linux support in the world today.




www.linuxgruven.com
888-443-9080



Open Windows Apps on Linux.

Visit www.netraverse.com for
free evaluation versions
of Win4Lin, a program that
lets you open and work
with Windows software on
a Linux OS.



Hey! LEGGO mysql

Installing, Configuring, and Using MySQL

Back in our September 2000 issue, we gave our Tuxie Editors' Choice Award for "Best Database" to MySQL. We've decided that it's time to feature some of the attributes that won MySQL that award. With easy installation and features such as database replication and integration with Perl, MySQL provides a scalable database solution for everyone from small developers to large database dependent applications.

First, if you don't have MySQL on your system already, you can download it from <http://www.mysql.com/downloads>. The latest, official, stable release (as of this writing) is 3.22.32. However, the newer 3.23.xx releases are now listed as "gamma" and are quite stable and recommended for new installations. The MySQL development team labels their releases conservatively. Therefore, in this article we will be working with the 3.23.30 release.

By Steve Suehring

Basic Installation

The MySQL Web site has pre-compiled MySQL versions available in both RPM and .tar.gz (tarball) packages. Like all free software, the source to MySQL is also available. We'll walk through the standard *configure*, *make*, and *install*, starting from the source distribution. MySQL should compile and install on any Linux distribution that has been released in the last few years, provided that you have *gcc* installed.

For our basic installation we will install MySQL into */usr/local*. We'll also create a *mysql* account and group that MySQL will run under so that we're not needlessly running it as root.

```
> groupadd mysql  
> useradd -g mysql mysql
```

Once the group and user are added, unpack the MySQL source.

```
> tar -zxvf mysql-3.23.30-gamma.tar.gz
```

cd into the newly created *mysql-3.23.30-gamma* directory and configure the installation.

```
> ./configure --prefix=/usr/local/mysql
```

There are many other options you can set at compile-time using additional command-line arguments to *configure*. They are all documented in the *INSTALL-SOURCE* file (which is a very detailed README).

If you have little memory in your system, you may need

to add the `--with-low-memory` option to *configure*. Otherwise, *gcc* may run out of memory during the compilation. In any case, once you're done configuring the installation, it's time to build MySQL. Execute the command:

```
> make
```

Depending on your hardware and system configuration, it will take anywhere from five minutes to an hour to build MySQL.

If the *make* fails for any reason, consult the `INSTALL-SOURCE` file for help. Most common compilation problems are listed with known solutions.

Assuming MySQL compiled without error, execute the following commands to complete the installation:

```
> su
> make install
```

After that's done, we'll need to run one of the scripts included with MySQL to create the default databases. For our purposes, we'll use the *mysql_install_db* script, which you can execute by issuing the command:

```
> scripts/mysql_install_db
```

Next, we need to change ownership of the newly created MySQL directory and files to the *mysql* user we created earlier.

```
> chown -R mysql.mysql /usr/local/mysql
```

Finally, it's time to start the database server:

```
> /usr/local/mysql/bin/safe_mysqld --user=mysql &
```

If something goes wrong, examine the contents of the error file in */usr/local/mysql/var*.

In case you are wondering, *safe_mysqld* is simply a shell script that launches the MySQL server process *mysqld*.

Before going any further, you should change the root password for the database.

```
> /usr/local/mysql/bin/mysqladmin -u root password new-password
```

MySQL does not rely on Linux's usernames and passwords. Each user needing access to the database must be listed in MySQL's internal user database.

You should now be able to log in to the MySQL database via the command:

```
> /usr/local/mysql/bin/mysql -u root -p mysql
```

As you might have guessed, the `-u` specifies the username you'd like to use when logging in. However, the `-p` tells the *mysql* command-line tool that you'd like it to prompt you for a password in this case. You *can* put the password on the command line, but there must be no space between the `-p` and the password. The *mysql* at the end of that command specifies which of the databases on your database server you'd like to connect to. Unless told otherwise, the *mysql* command-line tool assumes that you will be connecting to the MySQL server on the local machine.

So when you execute that command, you will be prompted for the password. Simply enter the *new-password* you just created.

Congratulations! You now have a fully functional MySQL database server running. It's now a good idea to set up the MySQL server to start at boot time and, more importantly, shutdown when the system is halted.

For Red Hat and other SYSV R4-type systems, simply copy the file *mysql-3.23.30-gamma/support-files/mysql.server* to */etc/init.d*. Next, symbolically link *S99mysql* into */etc/rc.d/rc3.d* and *S01mysql* to */etc/rc.d/rc0.d* as follows:

```
> ln -s /etc/init.d/mysql.server /etc/rc.d/rc3.d/S99mysql
> ln -s /etc/init.d/mysql.server /etc/rc.d/rc0.d/S01mysql
```

Remember to make */etc/rc.d/init.d/mysql.server* executable; it's not executable by default.

```
> chmod 755 /etc/rc.d/init.d/mysql.server
```

For systems that keep their *init* scripts in a different location, you'll need to consult your local documentation.

Server Configuration

Like many daemons, when *mysqld* starts up it looks for a configuration file. It will look for the file */etc/my.cnf* by default. So to make things run smoothly on your system, you'll need to build your own copy of */etc/my.cnf*. The MySQL distribution comes with four sample configuration files in the *support-files* directory. You should select one of those to use as the basis for your own. Which one to select depends on your hardware resources and whether MySQL is the primary service running on the system or if it must share resources with other services (such as Apache or Sendmail). The four included sample files are:

- *my-huge.cnf*: For systems with 1 GB or more RAM where MySQL is the primary service running on the machine.
- *my-large.cnf*: For any systems with 512 MB of RAM where MySQL is the primary service running on the machine.
- *my-medium.cnf*: For systems with less memory (32 MB – 128 MB) where MySQL is sharing resources with other services.
- *my-small.cnf*: For systems with less than 64 MB of memory where MySQL is sharing resources with other services and will not be utilized much.

For our purposes, we'll copy the medium file to *my.cnf* and edit it to suit our needs.

```
> cp my-medium.cnf my.cnf
```

The *Customizing my.cnf* Sidebar, pg. 40, contains a list of the configuration options contained in *my.cnf* and their defaults as they appear in the `[mysqld]` section. In this article, we are only going to focus on the options that you are likely to customize in a new installation. As you gain

Customizing my.cnf

Open *my.cnf* in your favorite editor and focus on the `[mysqld]` section, which is the section that deals with the MySQL daemon. The other sections are specific to various command-line utilities and are beyond the scope of this article.

```
key_buffer=16M
```

The `key_buffer` tells MySQL how much memory it is allowed to use for caching index data. Increasing this value can often yield dramatic improvements in MySQL performance. However, if this value is set too high it can adversely affect system performance (there won't be enough memory available for Linux and other processes). Our test system has 256 MB RAM, so we'll change this value to 64 M.

```
table_cache=128
```

The `table_cache` is the number of simultaneously open tables that MySQL can have. MySQL requires two file descriptors for each unique open table plus one for each concurrent client using a table. For our larger installation, we will probably need more than 64, so we'll increase this value to 128.

```
sort_buffer=512K
```

If you plan to frequently use `ORDER BY` or `GROUP BY` within your queries, you may benefit from increasing the size of `sort_buffer`. For our database installation, we will leave the sort buffer value at its default.

```
myisam_sort_buffer_size=8M
```

`myisam_sort_buffer_size` determines the size of the buffer used when repairing a table or creating an index. For our installation, the default value will be fine.

One important value missing from *my.cnf* is the variable that tells MySQL which user it should run as. Because we created a `mysql` user earlier, add the following line to the `[mysqld]` section of your *my.cnf* file.

```
user = mysql
```

experience with MySQL, you may find reasons to alter some of the options not listed here.

When you have finished editing the values that are listed in *Customizing my.cnf*, you'll need to copy the *my.cnf* file to `/etc`.

```
> cp my.cnf /etc/my.cnf
```

Now MySQL is configured to automatically startup and shutdown when your system does. When MySQL starts it will use the configuration options in `/etc/my.cnf`.

Security

Like any service, MySQL is only as secure as the machine on which it is installed. While keeping your system up-to-date with recent security fixes is important, you also need to understand the basics of MySQL's security system so that you can keep your data safe. MySQL provides extensive host, database, table, and user-level security, which can all be applied to make the data and the server more secure.

Network Security

If the host that runs the MySQL database is not located behind a firewall, one effective security measure is to simply change the default TCP port (3306) that MySQL uses to listen for connections. Edit the `/etc/my.cnf` file and change the value of `port` = from 3306 to an unused port on your system. Make note of the port you change the MySQL server to, as you will need to apply the change to any remote clients that need to connect to your server.

If your MySQL server is behind a firewall and will not need to answer requests from outside that firewall, it is best for the firewall to block connections to TCP port 3306.

Database-Level Security

Upon receiving a connection request, MySQL checks to see whether or not the user is allowed to connect to the database they've requested. This check involves not only the username and password that the client provides, but also the host from which they are attempting to connect. If the combination of user, password, and host is not allowed, MySQL simply closes the connection.

Table-Level Security

When a user attempts to execute an SQL query, MySQL checks the user's privileges to see if the user is allowed to perform the requested operation. Some users may only be able to read (`SELECT`) data, while others might be given full read and write access to the tables. In fact, MySQL

allows you to set privileges for each of the common SQL queries: `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. In addition to those, MySQL defines the privileges: `ALTER`, `CREATE`, `DROP`, `GRANT`, `FILE`, `INDEX`, `PROCESS`, `REFERENCES`, `RELOAD`, `SHUTDOWN`, and `USAGE`.

The rule of thumb with MySQL security is the same as that of any good system security policy: Only give users the access they need — no more. You will find that most users should only need `SELECT`, `INSERT`, `UPDATE`, and `DELETE` privileges.

Adding an Administrative User

Let's add a new user to the database. This user will be a database administrator who will be allowed to add databases and other users.

Log in to the MySQL server:

```
> /usr/local/mysql/bin/mysql -u root -p mysql

mysql> GRANT ALL PRIVILEGES ON *.* TO admin@localhost
IDENTIFIED BY 'password' WITH GRANT OPTION;

Query OK, 0 rows affected (0.03 sec)
```

Exit and then log in as the new administrative user to make sure it works:

```
mysql> exit
> /usr/local/mysql/bin/mysql -u admin -p mysql
```

Adding a Normal User

Now we'll add a normal (non-administrative) user who is allowed to add, update, and modify data stored in any table in the *users* database that we'll create in the next step.

```
mysql> GRANT SELECT,INSERT,DELETE,UPDATE ON users.*
TO dbuser@localhost IDENTIFIED BY 'password';

Query OK, 0 rows affected (0.01 sec)
```

It's useful to know that MySQL will let you assign permissions on objects (databases, tables) before they exist. This allows for planning security in advance of setting up your databases.

Creating a Sample Database

It's now time to create a database. We'll create a simple database called *users* that other services (like *RADIUS* or *POP3* servers) might use for authentication.

If you are still logged in, exit out of the MySQL client. Now you can create the database from the command-line as follows:

```
> /usr/local/mysql/bin/mysqladmin -u admin -p create users
```

Or, if you are already logged in to MySQL, you can use the `CREATE DATABASE` statement:

```
mysql> CREATE DATABASE users;

Query OK, 1 row affected (0.06 sec)
```

That's all there is to it. If you disconnected, you can reconnect to the new database with this command:

```
> /usr/local/mysql/bin/mysql -u admin -p users
```

Now we'll create a *passwd* table in which we can store basic user information:

```
CREATE TABLE passwd (
  username      varchar(8)      not null,
  crypt         varchar(128)    not null,
  realname      varchar(64)     not null,
  shell         varchar(20)     not null
);
```

Adding Data

Now we have a working administrative (*admin*) user, a new database (*users*) in which to store user information, and a user with minimal privileges (*dbuser*) to enter that information. Next, we'll insert a record in the *passwd* table using the MySQL `MD5()` function in an `INSERT` query. This will generate an encrypted password.

First, you will need to exit MySQL and log back in again as *dbuser*:

```
mysql> exit
> /usr/local/mysql/bin/mysql -u dbuser -p users
```

Now you can add a user to the *passwd* table in the *users* database and verify that the data is there with a `SELECT` query, as illustrated in *Figure One*.

For the complete guide to MySQL security, be sure to read the `GRANT` and `REVOKE` Entries in the MySQL Manual (see the *Resources* sidebar, pg. 42).

```
mysql> INSERT INTO passwd VALUES ('steve',md5('badpassword'), 'Steve Suehring','/bin/zsh');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM passwd;

+-----+-----+-----+-----+
| username | crypt | realname | shell |
+-----+-----+-----+-----+
| steve   | 31edaffbaba455bc30c52681ceb1ea9d | Steve Suehring | /bin/zsh |
+-----+-----+-----+-----+

1 row in set (0.00 sec)
```

Figure One

Resources

MySQL On-Line Manual

<http://www.mysql.com/documentation>

MySQL Download Site

<http://www.mysql.com/downloads>

GRANT and REVOKE Manual Entries

<http://www.mysql.com/doc/G/R/GRANT.html>

MySQL Related Perl Modules

<http://www.mysql.com/downloads/api-dbi.html>

MySQL Mailing Lists

<http://www.mysql.com/documentation/lists.html>

Perl

Having a database isn't terribly useful unless you can use it from your favorite programming language. One of the best features of MySQL is that we can use it from a wide variety of programming languages (Perl, Python, PHP, Java, C, C++, etc.).

We'll walk through the process of installing the necessary Perl modules that will allow you to use MySQL from a Perl script. Then we'll look at a simple Perl script that uses the modules to gather some information about your MySQL server.

Module Installation

If you haven't already downloaded the Perl modules from the MySQL Web site, you can find them at <http://www.mysql.com/downloads/api-dbi.html>. The modules are also available from any CPAN mirror site.

MySQL.com recommends downloading at least the following modules:

- Data-Dumper
- DBI
- Mysql-Mysql-modules

We also recommend that you download the Data-ShowTable module as well. These

modules should all be installed in the following order:

- Data-Dumper
- DBI
- Data-ShowTable
- Mysql-Mysql-modules

Once you've downloaded the modules, you'll need to be logged in as `root` to start the process of installing them. Like most Perl modules, these all follow the standard installation process, starting with the Data-Dumper module:

```
> tar -zxvf Data-Dumper-<version>.tar.gz
> cd Data-Dumper-<version>
> perl Makefile.PL
> make
> make test
> make install
```

Simply repeat the same steps for the DBI and Data-ShowTable modules.

The installation of the Mysql-Mysql-modules is slightly more involved. You'll need to make sure that your MySQL server is running before beginning the installation of Mysql-Mysql-modules.

The Makefile is interactive and will ask some questions about the install, as illustrated in *Figure Two*.

A Sample Script

Now it's time to build our first script. In *Figure Three* (pg. 44) we create a script that will connect to the `mysql` data-

```
> perl Makefile.PL
```

Which drivers do you want to install?

- 1) MySQL only
- 2) mSQL only (either of mSQL 1 or mSQL 2)
- 3) MySQL and mSQL (either of mSQL 1 or mSQL 2)
- 4) mSQL 1 and mSQL 2
- 5) MySQL, mSQL 1 and mSQL 2

Enter the appropriate number: 1

Do you want to install the MySQLPerl emulation? You might keep your old MySQL module (to be distinguished from DBD::mysql!) if you are concerned about compatibility to existing applications! [n]

Where is your MySQL installed? Please tell me the directory that contains the subdir 'include'. [/usr/local/mysql]

Which database should I use for testing the MySQL drivers? [test] mysql

On which host is database mysql running (hostname, ip address or host:port) [localhost]

User name for connecting to database mysql? [undef] root

Password for connecting to database mysql? [undef] <new-password>

Creating files for MySQL

Checking if your kit is complete...

Looks good

Warning: prerequisite Data::ShowTable 0 not found at (eval 14) line 228.

Using DBI 1.14 installed in /usr/lib/perl5/site_perl/5.005/i386-linux/auto/DBI

Writing Makefile for DBD::mysql

Figure Two

BE THE ONE



CHILLIWARE

S O F T W A R E

ACT ONE: SCENE ONE

FADE IN:

INTERIOR. BOARD ROOM Somewhere on planet Earth - DAY (EST)

Camera PANS slowly across a packed Board Room. The CEO of your company demands a complete business plan that will double the company's profits. The entire office pauses as a palpable tension floods the room. Suddenly all eyes are on you... the user of **CHILLIWARE's** full line of software solutions.

CLOSEUP - YOU - LIT from underneath

CEO

(Offscreen)

You are the one!

FADE TO PROMOTION WITH JUICY WAGE INCREASE.

ROLL CREDITS




iceSculptor™
Desktop Publisher



MOHAWK™
Apache™ Server Configurator



MENTOR™
Documentation Wizard



NEXXUS™
Contact Manager




XtremeICE™ OS

CHILLIWARE™
LINUX WITH ATTITUDE

Chilliware™, Inc. • 3550 Wilshire Blvd., Suite 1418, Los Angeles, CA 90010 • p:: 213.365.8700 f:: 213.365.1150 • www.chilliware.net

Linux™ is the trademark of Linus Torvalds. Apache™ is a trademark of the Apache Software Foundation. All other trademarks are the property of Chilliware™, Inc.

BY POPULAR DEMAND



BLACK



Happy Hacking Keyboard series

NEW ADDITION - BLACK MODEL

- ✓ Very Small Footprint (11.6" x 4.3")
- ✓ Full Size Key Pitch
- ✓ UNIX Friendly Key Layout
 - Sun/Type 5 keyboard based key layout
 - No Caps Lock Key
 - No Windows Key
 - CTRL Key is in Right Position

Happy Hacking Keyboard :
Multi-Platform model
Sun / SPARC, PC (PS2), and Mac (ADB)

Happy Hacking Keyboard LITE :
PC (PS2) model

www.pfuca.com
1-888-681-8878

base on the local MySQL server and list each of the tables it contains.

DBI is Perl's Database Interface. It abstracts the specifics of most database APIs (MySQL, Oracle, PostgreSQL, etc.) so that you can write scripts that are largely database independent with relatively little effort. To learn more about DBI, be sure to read the on-line documentation that was installed with the modules:

```
> perldoc DBI
```

Using the script in *Figure Three*, some basic SQL knowledge, and the DBI documentation, you can start building applications with Perl and MySQL in no time.

Linux's Next Killer App?

There's obviously much more to say about MySQL but this is all the space we have at the moment. Considering MySQL's ease of use and flexibility, it's easy to see why many people are calling it "Linux's next killer app." Look for more articles in upcoming issues.

Steve Suehring is a systems engineer at CoreComm. He can be reached at suehring@dangermen.com.

Figure Three

```
#!/usr/bin/perl -w

use strict;
use DBI;

my $user      = "root";
my $password  = "YOUR_PASSWORD";
my $database  = "mysql";

my $dbh = DBI->connect("dbi:mysql:$database",$user,
    $password) or die "Cannot connect to database: $!";

my $query = $dbh->prepare("SHOW TABLES");
$query->execute();

while (my $tablename = $query->fetchrow_array) {
    print "Database $database contains a table
        called $tablename\n";
}

$dbh->disconnect;

exit;
```

Here is what the output of this script looks like:

```
> perl dbtest
```

```
Database mysql contains a table called columns_priv
Database mysql contains a table called db
Database mysql contains a table called func
Database mysql contains a table called host
Database mysql contains a table called tables_priv
Database mysql contains a table called testac
Database mysql contains a table called testad
Database mysql contains a table called user
```


3:28 a.m.

The time one of your Linux web servers goes down.

The time we spot the problem.

The time we fix it.

The time you are peacefully dreaming and not worrying about your hosting.



rackspace.com 800.961.2888

Technology
Partners:



OUTSOURCING LINUX

EVALUATING YOUR LINUX-BASED MANAGED HOSTING OPTIONS

by Kevin Railsback

Linux has made impressive inroads into the conservative enterprise server room in the past year. Use of Linux has expanded outward into many companies that have traditionally thought of themselves as Windows or Unix shops. With backing from big names in the tech industry, and with the continued support of Linux geeks everywhere, Linux is fast becoming a mainstream corporate standard. In fact, Linux has become such a dominant force in the Internet server market that even businesses without in-house Linux expertise are looking at ways to integrate Linux into their Internet infrastructure.

At the same time, the costs and limitations of hosting a large Web site in-house are proving to be impractical for many companies. It can be prohibitively expensive to run a T1 or T3 line out to a company campus, even for a large corporation. Fortunately, there is an excellent solution available for these needs. In many cases, companies with and without in-house Linux expertise are turning to co-location and managed hosting facilities to house their servers.

WHAT IS MANAGED HOSTING?

Managed hosting offers all of the benefits of server co-location, including reduced bandwidth costs, guaranteed electrical power, physical security, and a controlled environment. However, managed hosting facilities take all of this one step further, providing full management and maintenance of the

server hardware and software. Many managed hosting companies will set up dedicated servers to meet your specifications (without requiring you to purchase hardware upfront).

Managed hosting can be particularly attractive to companies that don't have a large internal IT staff to support and manage a co-located server or ser-

ver farm. It is also often more difficult to establish and enforce a Service Level Agreement (SLA — which basically means that you get a guaranteed level of service for a fixed amount of money per month) with your internal IT department than with a managed hosting facility that is set up to provide high uptimes and comprehensive management.

WHY CHOOSE MANAGED HOSTING?

Managed hosting offers several advantages over traditional server co-location. With a managed hosting provider, maintenance of the hardware and software is the responsibility of the provider. Also, system backups and disaster recovery are typically covered under the terms of the service. SLAs guaranteeing uptimes from 99.9 percent to 99.999 percent (the famous "five nines") can be negotiated if you are willing to pay the extra cash for such obscene reliability. These are major hassles in a typical co-location setup. With traditional co-location,

you are responsible for all hardware maintenance and may have to handle your own system backups and software patching. Also, the co-location provider is only responsible for reliable power and network connectivity. Keeping the server in working order remains the responsibility of the owner (you). This is fine if you have the budget and staff to handle the upkeep, though still costly if you have unexpected downtime.

So, if you are looking for a quick and easy way to get up and running with a Linux-based server or cluster of servers and you don't want the day-to-day hassle of managing your ser-

vers, read on. We'll outline some of the basic offerings from several major providers of managed hosting services. This is not an exhaustive list of all managed hosting providers. Indeed, making a truly comprehensive list of managed hosting companies would be an impossible task, as many ISPs and existing co-location facilities are getting into the act.

Also, all of the providers we looked at offer multiple platforms (in hardware and operating system) under their managed hosting services. We concentrated on places that offer some type of Linux hosting; no Wintel-only providers were considered.

DELLHOST

www.dellhost.com

Worldwide, most hosting companies concentrate on the services that they provide while relying on other established companies to provide the actual hardware that powers those services. They buy servers, install their own custom software images, and configure them for customers. This is true in the shared hosting scenarios that a large number of companies offer and with the more business-focused managed hosting offerings that have become popular in the past few years.

Dell saw this trend developing and expanded their business offerings to include a full-fledged hosting arm. DellHost uses Dell's PowerEdge enterprise servers to provide customers with a wide range of hardware options for complete hosting solutions. They use Dell PowerApp servers for companies that need a dedicated basic Web server machine at a lower cost. All of them are running Linux, of course. At \$199 per month for a dedicated box, DellHost offers an excellent price versus performance ratio.

DellHost offers customers the option of handling their day-to-day administration themselves or having DellHost techs do the hardware and software maintenance, monitoring, and management. For \$299, you get

DELLHOST.COM

a PowerEdge 2450 with 866 MHz CPU, 128 MB RAM, and 9 GB SCSI drive. The server is RAID capable, in case you want to add additional drives. And if your company needs some serious performance, stock systems as beefy as a Quad Xeon 550 MHz with 2 GB of RAM and three 18 GB drives in a RAID 5 configuration are available. This isn't bad for \$1,249 per month with a 12-month contract.

If you're looking for solid managed hosting, a wide range of available hardware, and guaranteed uptime for a low cost, DellHost is just the ticket. While they don't offer platforms other than their own Dell servers, and center their services to a large extent on standard hosting, their managed hosting offering is among the best.

EXODUS

www.exodus.net

Exodus is one of the big boys of site hosting, server co-location, massive bandwidth, and top tier peering relationships. Many well-known Web sites are hosted on Exodus servers. In fact, they're the provider that VA Linux Systems uses to host sites in their Open

Source Developers Network. Even that company in Redmond uses Exodus for hosting Internet and internal servers.

Exodus provides managed hosting for a wide range of hardware and software platforms. Customers can rent preconfigured servers from a variety of vendors, including Sun, Dell, and Compaq. With their Managed Monitoring Services, a comprehensive managed hosting service can easily be set up. Exodus is all about choice and they offer various levels of service at different prices.

Exodus manages a worldwide network of Internet Data Centers in North America, Europe, and Asia. Through this widely dispersed network, vital data can be quickly replicated worldwide and then served up to end users from the nearest link. This not only gives the end user the fastest possible experience, but it also saves bandwidth by serving up information locally.

While managed hosting isn't Exodus' only business focus, they have the infrastructure and technical know-how to do an overall excellent job of supporting all sizes of Internet sites. Of course, their services are on the pricey end. However, you definitely get what you pay for. Exodus is a good choice for those sites that demand the



EXODUS

best in peering relationships and solid uptimes — and have a large enough budget to afford it.

HOSTOPIA

www.hostopia.com

If you can ignore the flying pig on the main page of their Web site, Hostopia might be your managed hosting utopia. Their approach is a bit different from the others we looked at. Hostopia is focused on providing services to companies that want to be hosting providers themselves. In other words, you can set yourself up as a full-blown hosting provider and have Hostopia

invisibly do the grunt work. It'll appear to your customers that you're providing the hardware and software services.

Hostopia provides a totally outsourced Web, e-mail, and e-commerce hosting solution. This is perfect for ISPs and ASPs who want to expand their services to include managed hosting for their customers. They offer cross-platform integration in their server farms, so customers are able to choose between BSD, Linux, and Windows services.

Hostopia's solution offers some unique advantages. For in-

stance, all administration and configuration is done through a customized Web interface, both for its direct customers and the end users of those customers. In addition, as a direct customer of Hostopia, you can manage your own customers' sites and account information through their w3Control Panel.

Hostopia isn't for everyone, but if you're an existing ISP or ASP and want to easily resell managed hosting services without investing in a huge server farm, they are definitely worth a look. On the other hand, if you're simply an end user or



COMPANY NAME	PURCHASE/LEASE EQUIPMENT	BANDWIDTH AVAILABILITY	SERVICES SUPPORTED	HARDWARE RANGE
DellHost	Purchase/lease	Excellent peering	Systems, Hosting, Security Networking, Application services	Dell
Hostopia	Lease	Excellent peering	Managed servers, Hosting	Intel, others
HostPro	Purchase/lease/ customer owned	Excellent peering	Systems, Colocation, Hosting, Dedicated managed servers, Connectivity, Application services	Intel, SPARC, Cobalt RaQ
Hurricane Electric	Purchase/lease/ customer owned	Excellent peering	Co-location, Managed hosting	Intel, SPARC, others
ManageNet	Purchase/lease/ customer owned	Good peering	Managed Web hosting, Database, E-mail, Internet access, LAN	Intel, SPARC, others
NonStopNet	Purchase/lease/ customer owned	Excellent peering	Managed Web hosting, Storage, Database, Networking	Intel, SPARC, others
RackSpace	Purchase/lease/ customer owned	Excellent peering	Managed Web hosting, Monitoring, Networking, Software services, Security, Database, Backup	Intel, SPARC, others
ValueWeb	Purchase/lease/ customer owned	Good peering	Dedicated servers, Commerce hosting	Intel, Cobalt, SPARC, others
Verio	Purchase/lease/ customer owned	Excellent peering	System admin, Backup, Web hosting, Site design, Security, ASP solutions, Security, Load balancing	Intel, SPARC, others
Exodus	Purchase/lease/ customer owned	Excellent peering	Systems, Networking, Content distribution/caching, Security, Storage, Performance solutions, Application services	Sun, Compaq, Dell, others

a small business who only wants basic managed hosting, Hostopia's services aren't for you.

HOSTPRO

www.hostpro.com

HostPro is the perfect example of a diverse hosting provider. They offer traditional co-location, dedicated servers with managed hosting, individual site hosting, application hosting, and even home and business Internet access services. This makes them an excellent choice for non-technology companies looking for one source for all of their Web needs.

HostPro's wide platform support and varied services makes them a great pick for tech companies also. They offer dedicated hosting on Linux, Cobalt RaQ 3 and RaQ 4 servers, Windows NT, Windows 2000, and Solaris platforms. All of these platforms can be set up as managed or unmanaged, allowing companies to pick and choose how they want their hosting servers set up.

HostPro offers a comprehensive site control panel that allows the customer's administrator (or others in the organization) to set up sites, configure options, set up e-mail accounts, cus-

HOSTPRO®

tomize site security, and view site statistics. Basically, any-

thing that you would expect an experienced admin to handle can be done by a non-geek through a control panel without command-line tweaking. Of course, if command-line tweaking is your thing, it's available too.

A standard HostPro managed Linux server (the DH6) comes with a 650 MHz Pentium III CPU, 256 MB of RAM, 9 GB SCSI disk, and 50 GB of monthly bandwidth for \$1,045 per month. This can be scaled up for those who need even more horsepower. And if you're looking for the ultimate in

OSes SUPPORTED	CLUSTERING	ROOT ACCESS	SLAs	PRIMARY AUDIENCE	PRICE RANGE	URL
Linux, Windows NT/2000	Yes	Yes	Yes, 99.9% uptime guarantee	Large enterprise customers, mid-sized businesses, Internet businesses	\$199 per month for basic Linux PowerApp server, \$299 per month for basic Linux PowerEdge server	www.dellhost.com
Linux, BSD, Windows NT/2000	Yes	No	Yes, 99.99%	ISPs, ASPs	\$9.99 to \$69.99 per account	www.hostopia.com
Linux, Solaris, Windows NT/2000	Yes	Yes	Yes, wide range	Small to large businesses, individuals	\$1,045 per month and up for managed Linux servers, \$495 for unmanaged	www.hostpro.com
Linux, Windows NT/2000, others	Yes	Yes	Yes	Small to mid-sized businesses, individuals	Unpublished	www.he.net
Linux, Unix, Windows NT/2000, Novell	Yes	Yes	Yes	Small to large businesses	Unpublished	www.managenet.net
Linux, Unix, Windows NT/2000, others	Yes	Yes	Yes	Mid-sized to large businesses	Unpublished	www.nonstopnet.com
Linux, Unix, Windows NT/2000, others	Yes	Yes	Yes, 99.999%	Individuals, Small to large businesses	\$380.80 for baseline Linux server	www.rackspace.com
Linux, Unix, Windows NT/2000, others	Yes	Yes	Yes	Individuals, Small to mid-sized businesses	\$295 per month for basic VA Linux server	www.valueweb.net
Linux, Windows NT/2000, others	Yes	Yes	Yes, 99.9% and up	Individuals, Small to large businesses	\$375 for basic Linux dedicated server	www.verio.com
Linux, Unix, Windows NT/2000, others	Yes	Yes	Yes, wide range	Large enterprise customers, mid-sized businesses, Internet businesses	Unpublished	www.exodus.net

performance and scalability, HostPro also has managed clustering systems. A pair of clustered DH6 servers with a dedicated back-end storage server runs \$3,250 per month and ensures maximum uptime for mission-critical applications. A cluster of four DH6 servers with a back-end storage server is available at \$6,050 per month, with even higher-end custom options for your super-duper, 24/7, no-downtime needs.

HostPro offers a wide range of managed hosting options and services, making them a good choice for a wide range of users. They have low-end servers available for the cost conscious and massive, geographically distributed clustered servers for those with the most demanding needs (and large budgets).

HURRICANE ELECTRIC

www.he.net

Hurricane Electric sounds like some weird weather research company or power generating station. But they are a well-known Web hosting provider. They've made the leap from hosting and server co-location into the managed hosting arena, and their services match many larger providers.

Dedicated servers with a gigabit backbone make Hurricane Electric a good choice for bandwidth-intensive customers. They handle all aspects of getting your site up and running, from ordering and configuring your server to taking care of the long-term system management. They have spare hardware on site in case of hardware failures, technicians ready to put out fires 24/7, and massive bandwidth to keep your site pumping even under Slashdotted loads.

Unfortunately, only their self-service Web hosting services are available for instant setup from the Web. However, if you're not in a last-minute rush to get your site on

the Net, and are on a tight budget, Hurricane Electric may be the service you need to get started.

Hurricane Electric's services are geared primarily towards individuals and small to mid-sized businesses, although they have the bandwidth and peering relationships to support larger businesses. However, if you're in need of high-end servers and rock-solid SLAs, a small site like Hurricane Electric might not be the best fit. For the budget conscious managed hosting shopper however, Hurricane Electric may be just what you're looking for. Now if they could get rid of that power shortage in California...

MANAGENET

www.managenet.net

As their name implies, ManageNet is a company dedicated to managed hosting and services. They offer a wide range of managed hosting options, including managed Web hosting, database servers, and managed global e-mail services. ManageNet also offers Internet access and managed LANs and data centers for companies who want to outsource all of their IT resources.

ManageNet's local hosting services are a unique approach to linking your business to the Web. ManageNet will connect your existing IT infrastructure to key partners, customers, and remote sites, effectively doing managed hosting of your internal network as well as remote managed Web hosting. Of course, you can pick and choose how far you want to take your company's relationship with ManageNet. If you're looking for standard managed Web and application hosting, they can take care of that with ease.

Hardware and software support is broad with ManageNet, including Linux (of course), many flavors of Unix, Windows NT, and even NetWare services. They offer a range of pricing and service levels, which allows companies to start off with



basic managed Web hosting services and then grow into their more in-depth network management and back-end hosting. ManageNet also offer services without long-term contracts, an appealing option for those just starting off on the Web.

ManageNet isn't for everyone. If all you'll ever need is basic managed Web and e-mail hosting, many of their capabilities will be lost on you. However, if you hope to grow your managed hosting relationship into a broader overall IT outsourcing and management arrangement, ManageNet may be the perfect combination of services and costs.

NONSTOPNET

www.nonstopnet.net

NonStopNet is a managed hosting provider with one goal — redundancy at every corner. They offer varying ranges of managed hosting, with geographically distributed data centers and a highly scalable infrastructure for businesses that need their hosting provider to grow with them. NonStopNet offers managed Web hosting, storage, databases, and even entire managed networks for customers.

Their service isn't limited to providing the managed hosting that your company needs. Instead, they work with you from the needs analysis stage through network design to implementation. It doesn't stop there. Part of NonStopNet's service is an ongoing relationship in which they help your company reevaluate current needs and tweak the service accordingly.

NonStopNet can provide high-performance data storage services for your business as well. Replication between sites ensures availability, and managed clustering ensures data integrity. Of course, these high-end services come at a cost, but if your busi-



HURRICANE ELECTRIC
INTERNET SERVICES

eLinux

Linux Technology Solutions

Cobalt RaQ™ 3
Server Appliance
\$999

only
Save \$100!

Save up to \$200 on Cobalt Servers!

eLinux has the lowest
prices on Cobalt Servers
guaranteed!

Call today!

Cobalt RaQ™ 4
Server Appliance

only **\$3,099**

Save \$200!

Cobalt RaQ™ 4i
Server Appliance

only **\$2,549**

Save \$200!

Cobalt RaQ™ XT
Server Appliance

Call for best price!

eLinux—Your total Linux solution!

Dedicated Account Executives

Multi-Vendor Solutions

Custom Configurations

Online Resources

Pre and Post Sales Support

True Blue Diamond Partner



**COBALT
NETWORKS**

Mention source code LXCBLT

Call toll free 1-877-39-LINUX

Fastest Overnight Delivery! Order today—Get it tomorrow! Call toll-free 1-887-395-4689

ness depends on these mission-critical services, the costs of not having this support can be devastating.

Massive redundancy and high availability come at a higher overall cost. NonStopNet is probably not the best choice for small businesses or individuals who want basic managed hosting, although they do provide those services. Instead, NonStopNet is a better choice for mid-sized to large businesses that need high availability and plenty of room to grow their hosting needs.



NonStopNet

managed servers are, etc.

Once the operating system is chosen, a number of preconfigured options are available. Each can be customized. If you choose a stock configuration, Rackspace can get your site up and running in an hour!

Rackspace is a popular choice for those just getting started in the managed hosting arena, and for good reason. It's easy to set up, customizing your server is simple, and costs can be matched to your budget. If you're looking for an all-around strong managed hosting service and you don't want to spend an arm and a leg to get started, you should give Rackspace a close look.

who need to get their commerce site going without the expense and hassle of doing it themselves.

ValueWeb's dedicated servers offer enormous variety. They offer Intel-branded server hardware, VA Linux servers, and even Cobalt RaQ systems. ValueWeb has a wide range of hardware configurations that can be customized to meet your business' exact needs. If, on the other hand, your company already has their own servers, ValueWeb will also co-locate your servers at their data center.

ValueWeb™

This means that you can start with shared or merchant hosting and then grow into a dedicated server or fully customized company-owned solution without needing to change providers. ValueWeb doesn't offer the in-depth managed hosting of the other sites we discuss in this piece, but their co-location and dedicated server solutions are worth a look for those that don't need someone managing their Linux boxes.

RACKSPACE
www.rackspace.com

When many people think of managed hosting the first company that comes to mind is Rackspace. They have made a name for themselves by allowing the easy and quick setup and configuration of a wide range of dedicated servers at reasonable prices. Rackspace focuses exclusively on the managed hosting market, and it shows in the way that their site and service are set up.

With Rackspace's "Price a Server" pages, customers can see (and modify) exactly what they get for their money. For example, each server component, from CPU speed to operating system, can be modified and the price can be recalculated for monthly managed hosting. There is no need to purchase your server up front, as its cost is already factored into the monthly hosting fee.

The ease with which you can tweak server options makes setting up a Rackspace server fun. The first option you choose is your operating system. Linux is the first choice in the list (of course), but you can choose from Solaris, FreeBSD, Cobalt RaQ, or even Windows NT/2000. For non-technical users, they even offer assistance with figuring out the basics, such as how to pick an OS, what

VALUEWEB
www.valueweb.com

ValueWeb is a hybrid — more than a co-location facility but less than a full managed hosting provider. If you are looking for a company that can provide dedicated servers, 99.9 percent uptime, and a money back guarantee, and are willing to manage your Linux server yourself, then ValueWeb may be a good option.

ValueWeb is less expensive than many of the other companies we list in this article, but their service does not extend to the day-to-day management of your servers. They provide the common shared hosting options that many other sites offer and give growing or larger companies the option of having their own dedicated, custom-configured server.

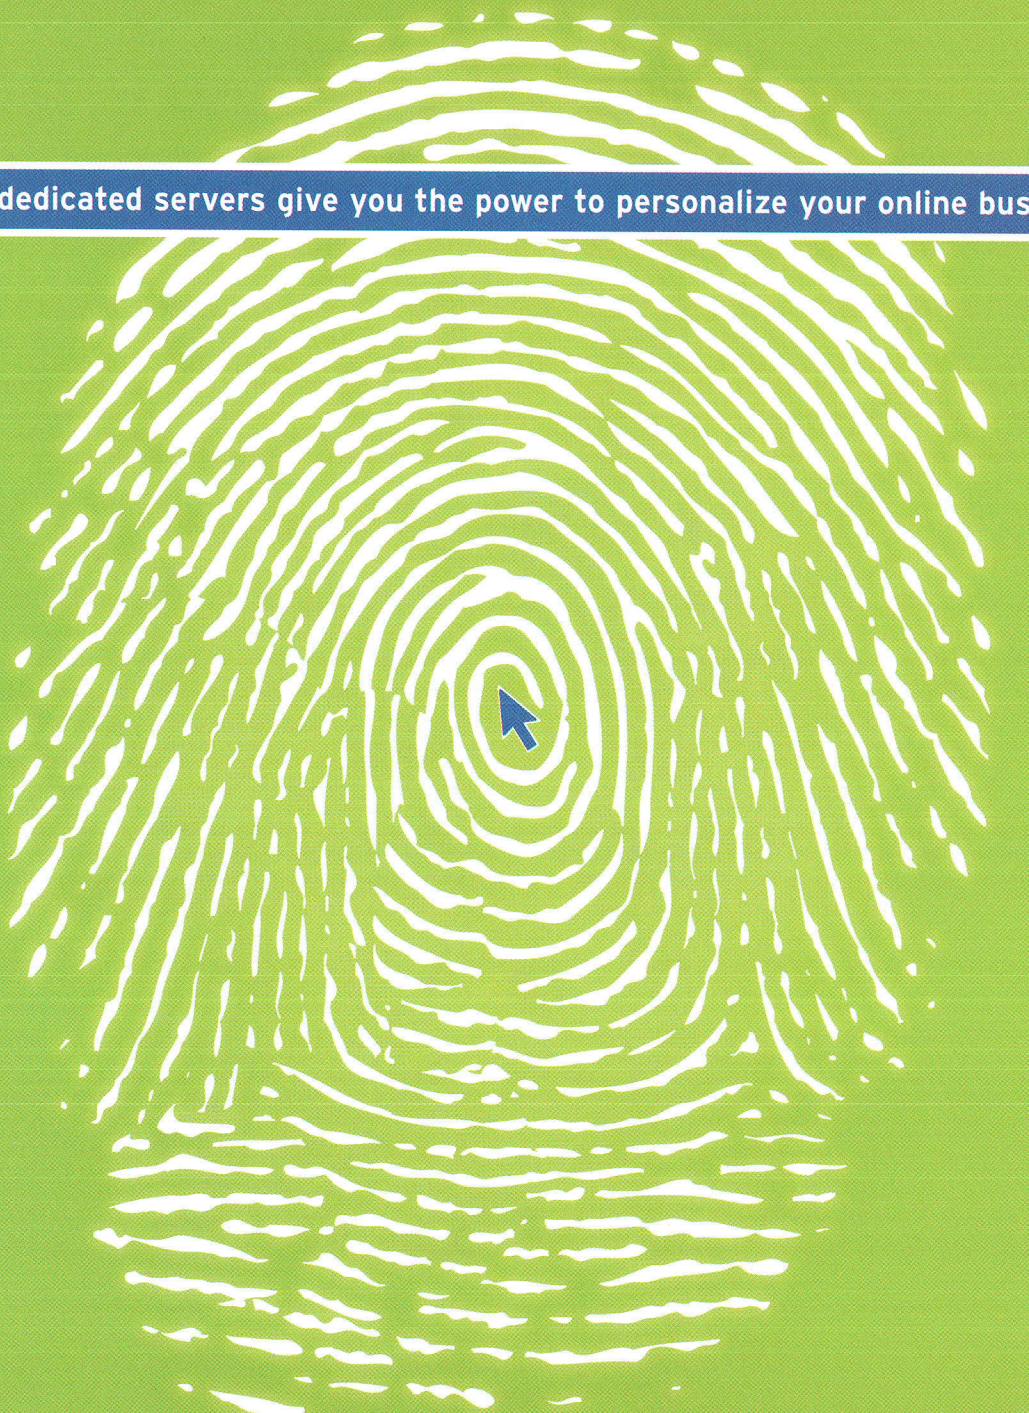
ValueWeb can do more than set up a simple server however. They offer merchant hosting options that allow new Web-based businesses to get up and running quickly with a full suite of commerce tools. While not as powerful in the long run as a totally custom server that your company must design and manage, their merchant hosting is a perfect choice for those

VERIO
www.verio.com

Verio is another big name in Web hosting services. They offer dedicated servers, high bandwidth, rock solid facilities, and of course, managed hosting of your dedicated servers. Similar to HostPro and ManageNet, Verio can provide a full range of services from basic Internet access for your business to managed clusters of high-end servers.

The managed services that Verio offers are divided into different categories. In this way, your company can choose which management pieces it needs handled by an outside source while you can take care of some aspects yourself, if you so desire. The categories are: system administration, backup services, security/firewall services, server monitoring services, and server load balancing. If you choose


rackspace
MANAGED HOSTING



Verio dedicated servers give you the power to personalize your online business.


Verio Dedicated Servers. The power of a fully customized website is just a dedicated server away. Intranet software, e-mail, message boards, e-commerce applications and streaming media are all possible when the server is exclusively yours. For a monthly fee, Verio will manage the server hardware, operation and maintenance. As a hosting pioneer, our staff, partners and Tier One network offer maximum speed, scalability and security. And that's the kind of power that lets you put your own unique stamp on your e-business.

VERIO
the new world of **business**

Toll Free 877.624.7894 e-mail: dedicatedserver@verio.net www.dedicatedserver.com

W I N D O W S 2 0 0 0 / N T S U N L I N U X C O B A L T

© 2001 Verio Inc. Verio and The New World of Business are trademarks of Verio Inc. All other referenced product names are trademarks of their respective owners.



Copyright 2001, Terra Soft Solutions, Inc.

2.0

YDL 2.0 for PowerPC

by Terra Soft Solutions, Inc.

*Featuring kernel 2.4, KDE 2, XFree86 4
and a graphical installer like none other.*

www.yellowdoglinux.com
Redefining your PowerPC experience.

to only subscribe
to some of these
services you can
still have help

with any of the others by simply paying an hourly or
per-job fee for system administration support.

Similar to the other facilities we discussed, Verio
gives you full root access to your servers. They also
have peering relationships with 15 Internet backbone
links, ensuring that your site will be accessible even
in the event of an Internet hub problem.

Verio is definitely a provider aimed at the business
community. They offer services to help your site grow
and prosper, such as Web site promotion services and
affiliate programs. They can handle the whole range
of Internet services that your company needs, from
registering your domain name to propagating your
storage needs and data between data centers.

If you're an individual looking for an easy to set up
place to host a simple Web page, Verio can do the job
for as little as \$24.95 per month. Of course, if you're
in need of managed hosting for a dedicated server,
you should look to sites like Rackspace or Hostpro.
Verio's higher-end services are well designed for
businesses that need to get their e-business going.


WHICH IS RIGHT FOR ME?

As many managed hosting companies as there are in
this article, this list is by no means a comprehensive
one. Any of these companies can provide basic Lin-
ux-based managed hosting, but the one you should
choose depends heavily on your budget and needs.

Large enterprises needing geographical server dis-
tribution and high-availability clustering should look
to providers such as Exodus, Verio, or NonStopNet for
their managed hosting. Small and mid-sized busi-
nesses that need a more cost-effective solution, while
maintaining a high level of service, should consider
Rackspace, DellHost, HostPro, NonStopNet, Hurricane
Electric, or ManageNet. ISPs and ASPs who are look-
ing to resell hosting services will find what they need
with Hostopia, while individuals should consider
less costly alternatives such as Rackspace's low-end
offerings, Hurricane Electric, or HostPro.

When you're looking for a managed hosting pro-
vider, the key is to make sure you know exactly what
you need, what you're willing to pay, and what your
future plans are. There are literally hundreds of host-
ing providers out there, and each offers a unique com-
bination of services and features (at a wide range of
cost); you should choose carefully.

*Kevin Railsback is the west coast technical director for
the InfoWorld Test Center. He can be reached at kevin_
railsback@infoworld.com.*



Data Storage and
Access Solutions for:

Telecommunications

Research

Education

Health Care

eGovernment

Financial

OTG SOFTWARE. TAKING THE WORLD BY STORAGE.

New for Linux/UNIX system users.
Infinite storage capacity for critical data.

OTG Software's DISKxtender™ takes Linux and UNIX systems to a higher level of storage performance. That means more enterprise data can be stored, accessed and managed more efficiently than ever. You'll have access from desktops, intranets, even the Web of business data. Today's e-world demands greater performance from your storage capacity. For Linux and UNIX users, the solution now includes DISKxtender™.

INFINITE STORAGE. ULTIMATE ACCESS.

www.otg.com

T 800 324 4222 World Headquarters
T 301 897 1400 6701 Democracy Blvd., 8th Floor
NASDAQ: OTGS Bethesda, MD 20817





The feeling you've found where having control is

Introducing **FREEDOM**. More power for less, for once.

It's called HostPro Freedom for a reason. For the first time ever, HostPro gives you more control than servers costing three times as much. We're the only host to offer you four unique features for Unix platforms. First, you have complete control over applications. HostPro Freedom lets you run just about any software you want, or even change the Web server. Second, customization is no problem, since you're running open source FreeBSD and have access to a port library of 4000 tested apps. Unlimited shell accounts allow you to create and control multiple

users. Third, you get multiple domain hosting for just \$149.95 a month, with up to 40 domains hosted at no additional monthly charge. And finally, HostPro Freedom's Virtual Integrity technology adds unique stability to your site. Each account acts as though it has its own hard drive—a virtual disk, as we like to call it. That means your shared server stays stable even when other users run their own applications.

Impressive? We think so. No one else offers all this freedom. Call us for more information, or go to HostPro.com. And then do anything you want.



**the one place
complete Web-hosting
ridiculously affordable.**

Specifications

	FREEDOM 200	FREEDOM 300	FREEDOM 400x	FREEDOM 500x
Price	\$59.95	\$99.95	\$149.95	\$249.95
Disk Space	300MB	400MB	500MB	1GB
Bandwidth	10GB	20GB	30GB	40GB
Email Accounts	Unlimited	Unlimited	Unlimited	Unlimited
Shell Access	Yes	Yes	Yes	Yes
Root Access	No	No	Yes	Yes
FTP/Shell Users	1	10	Unlimited	Unlimited
Multiple Domain Hosting	10 Domains	20 Domains	40 Domains	Unlimited
Shopping Cart	Miva Merchant	Miva Merchant	Miva Merchant	Miva Merchant
SSL Certificate	Yes	Yes	Yes	Yes

(visit our website for more)

HOSTPRO®

Try us on for size.

1-800-862-8917

hostpro.com -- enter code: BIGGUY4
BIGGUY4@hostpro.com

© 2001 HostPro



SPARKING

The Embedded Revolution



Lineo's Bryan Sparks Wants Linux to Go Anywhere

Does Linux really spell the end of proprietary embedded operating systems? Lineo CEO Bryan Sparks certainly thinks so — and he's doing everything he can to make it happen.

by Robert McMillan

When all is said and done, Linux may have the greatest impact where it is the hardest to see — in the phones and PDAs and routers that, more and more, are being powered by the Penguin. One of the people at the forefront of the embedded Linux phenomenon is Bryan Sparks, CEO of Lindon, UT-based Lineo. After his Linux company Caldera bought DR-DOS from Novell in 1996, Sparks began building a new line of business selling software and services to the embedded market. Before long, Sparks realized that it was Linux and not DR-DOS that held the most potential. Two years later, the embedded business was spun off from Caldera, and Lineo was born. Lineo filed for an IPO in the fall of 2000. Due to unfortunate market conditions, they were forced to call it off in January 2001. Despite this, when Sparks recently sat down to dinner with *Linux Magazine's* Publisher Adam Goodman and Contributing Editor Robert McMillan, he was highly optimistic about the future of both embedded Linux and his own company.


LINUX MAGAZINE: How did Lineo come to be spun out of Caldera?

BRIAN SPARKS: Two years ago we weren't even called Lineo. We were running under the awkward name "Caldera Thin Clients." We were building these TV set-tops, software solutions. We had a browser that we were writing just for TV output, and it was not even Linux-based. It was DOS-based. At that time, there was this guy building a DSL router who came to us thinking that we would do a DOS-based router deal. We said, "You don't want to do that. Tell you what, we'll help you with Linux." So we did a deal with him and put embedded Linux on his device. Then we got approached to put Linux on another device by a big device manufacture. We did a flash driver for them and that's when we said, "Gosh, what are we doing? You know what? We need to be doing embedded Linux." So in early 1999, we started talking to Motorola and a bunch of other vendors. Motorola caught on first and said, "We're thinking about doing Linux," and they showed interest and helped us define what products we should build. So, that's how we got here.

LM: In the last two years, has the embedded Linux market matured in ways that surprised you?

BS: Yes. Two years ago I thought I would find all our

success in areas where Linux was strong: networking, routing, firewalls, gateways, VPN [Virtual Private Networks], those kinds of things. The big surprise for us has been the demand for Linux for consumer electronics. It was not what we had anticipated. Linux alone doesn't have an intrinsic GUI for embedded systems. You read about X-based applications. Well, X doesn't work on embedded devices. It's a pig no matter how small you try to make it. It just doesn't work. However, even without an intrinsic GUI, there was still interest for Linux in embedded, consumer electronics that have graphical user interfaces, simple as they might be.



LM: Why are these consumer electronics developers interested in Linux?

BS: I think it's because we offer a richer technology base and a richer technology roadmap than the traditional RTOS [Real Time Operating System] vendors offer. It's not that it's open source, though that's interesting too because it gives the OEM (Original Equipment Manufacturer) a little more freedom than they would have if they were tied to someone like Wind River, who charges them for fixes to bugs that they find.

Certainly cost is involved, because Linux is perceived to be cheaper and indeed is. But I really think that it's a richer technology roadmap. I can do more, faster than an RTOS vendor, sometimes on an order of magnitude. We matched VxWorks' [Wind River's embedded operating system] functionality in our first release of Embeddix

[Lineo's embedded Linux distribution]. With our second release — man, we're way ahead. I mean, we're doing stuff I think they're just now starting to talk about.

LM: Is that because there is a community of developers that are contributing to the same core technology?

BS: Of course — and there's my ability to attract quality engineers and make them effective and productive sooner.

LM: So what is the difference between you and the other embedded Linux vendors? For example, you seem to approach real-time Linux in a different fashion from MontaVista.

BS: In order to provide hard real-time guaranteed determinism, there are several things we have to do. If you're an application on Red Hat and you want hard real time on Red Hat 6.2, you can't get that, because you are at the whim of the scheduler that schedules you. You can be a high-priority task, but if there are five other high-priority tasks, you're dead. So, we've had to introduce a layer of determinism that architecturally sits underneath the Linux kernel. Computer science academics call it an executive. It's a real-time executive that runs Linux as a task.

LM: Like a microkernel?

BS: Like a microkernel, yes. But, you see microkernels — Mach or Chorus or some of the others — those were micro only on paper. In actual implementation they were macro. Man, they were huge! Mach was just massive. And with Mach and Chorus you ran operating system personalities on top. We're not that full featured but we do have this little layer underneath, this run-time executive. And it's not an OS; we don't provide all those services. We expose an API and we have a whole developer kit that allows tasks to register their desire for

real-time services. That's how we do it. That's hard real time. We also do soft real time, which is really just having high priority tasks. That's what Monte Vista does. Monte Vista went in and modified the Linux scheduler. We did the same thing to essentially create a higher priority granularity in the scheduler.

LM: It seems like there are a variety of different forks of the embedded Linux kernel out there.

BS: What we don't want is... OK, we splintered Linux but not on purpose. We basically put Linux on devices that the community doesn't have momentum on — like an ARM chip. The community, in general, has made no effort to put Linux on the ARM chip. In that sense we're kind of a derivative. We don't mean to be. We try to give back but we're not normally picked up by the community because, again, there's just not a lot of interest. So in some areas we do branch. In the real-time space we're probably branching a bit too because, in general, the community doesn't really care.

LM: I've wondered about this. Everyone talks about the embedded market as an important space for Linux, but I'm not sure it really means anything to the community running Linux on their desktops and servers. Do you think it will?

BS: I don't know if it will or not. You think of the Linux community as a great big mass of people, right? So they all congregate around what the mass feels is interesting. Servers are interesting. For example, with Linux 2.4 there's a great big push for a better SMP. Those kinds of things don't really play in the embedded space. We play on ARM chips and Hitachi chips — things that are not interesting to the community at large. You can't buy a Hitachi-based PC. So I'm not sure there really is a connection to the community.

LM: Are you saying you don't really feel connected to the Linus Torvalds and Alan Coxes of the world?

BS: Oh, we do. We make changes and all of that. However, we do things that make the community in general go, "Boy, those are weird people over there."

LM: Do you think that it's inevitable that there will be many kernel forks with all these weird changes?

BS: We move to some different — I don't want to say weird, but certainly divergent from the community — platforms. I'll give you a concrete example. We run on devices that have chips that don't have an onboard hardware MMU [Memory Management Unit]. So there isn't an MMU, and so you have to modify the Linux kernel, which makes substantial use of hardware MMUs. For example, the `exec` system call is copy on write; that's an MMU function. We have to write our own `exec` for MMU-less chips. That's a fork.

LM: So do you have to consciously find ways to work with other embedded Linux vendors to minimize forking and try to get the most out of the work that you're all doing?

BS: Well, you know what? That was the original intent of the Embedded Linux Consortium. Perhaps we did that wrong, but that was what we wanted to accomplish. It's not going to do that unfortunately.

LM: Why not?

BS: I mean, Wind River is a member. They don't even talk Linux anymore. I don't even know what they are doing. The notion was, "Let's create conventions since we're in virgin territory for the Linux community." One obvious thing is how to launch apps on embedded devices that don't have command-line prompts? Let's just come up with a convention. Nobody

gets a competitive advantage out of it. Let's just come up with one we all agree on. That was what I wanted from the Embedded Linux Consortium. It didn't do that. It quickly turned into, "Well, I want my time," and, "Oh, we should be marketing," and ...oh man!

So now we're just kind of doing our own thing. I think we'd all agree that nobody sees a competitive advantage in the way we launch applications.

LM: Is the name Linux going to really mean anything in the long run if you guys have all these little forks and your competitors do things in different ways?

BS: It may not matter. Do we care if it's Linux? Well yeah, I don't want to shoot myself in the foot. My biggest differentiation right now is that I'm Linux. Linux is a means to an end for us, so it would be silly to drop Linux and do BSD. However, we are not a company that believes everything we build must be open source.



We don't challenge the GPL. In fact, we're working with the Free Software Foundation to actually define it a little better — but we're going to add proprietary value and we think we should get paid for our value-add. If you don't want it, that's fine. Build it yourself. That's capitalism in its purest sense.

So it's not that we're just forking Linux. I don't want you to get the impression that we are. We're just taking Linux into areas that the community in general isn't going into. I don't call that a fork. You might, but we don't really care if the community picks up the things we add because we're not publishing source to them.

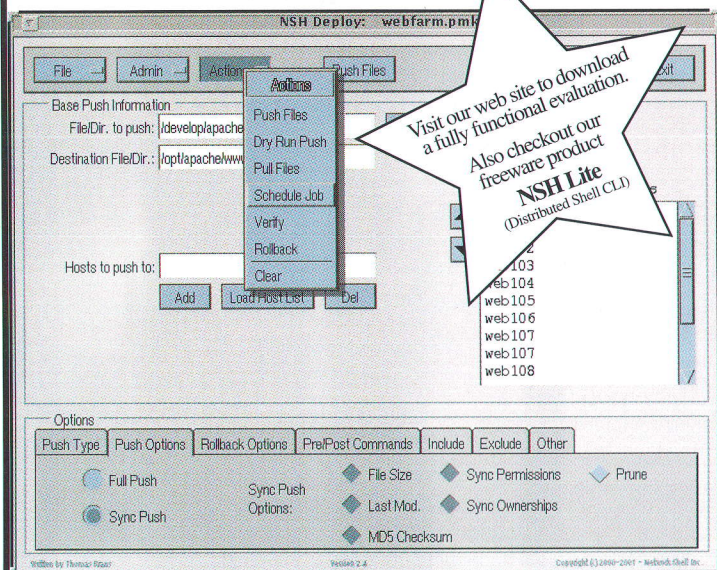
LM: Then why did you choose Linux?

BS: The paradigm shift — If I'd done just another RTOS, I would have had no chance of knocking Wind River out.

LM: So you're riding the momentum essentially?



On your Web Farm ... are all of your animals equal ?



When it comes to the management and deployment of your e-business applications and data, consistency is the key. With NSH Deploy, not only will you be up and running quickly, you can keep your web farm from becoming an animal farm!

NSH Deploy Features/Benefits

- Massively parallel pushing for minimal deployment time
- Job scheduling for automated deployment
- Pull files as well as push files for extended functionality
- ACLs and encryption for maximum security
- Independent verification for *anytime* consistency checking

For a fully functional evaluation please visit or contact

WWW.NETWORKSHELL.COM

1 888 I WANT NSH or 650 548-0674

BS: Well, we're riding the momentum and taking advantage of the inherent qualities that differentiate us from our largest competitors. How do you take your competitor's greatest strength and turn it into his weakness? You can't change his strength.

Wind River's strength is VxWorks. They can't just swap that out, because \$400 million a year in revenue would simply disappear. So you ask, "How can we take that out?" Well, you know what? We commoditize it. Wind River talks about raising per-unit value in time; they tell their investors that. I say we're here to commoditize system software for embedded devices. That hurts them and gives me opportunity.

I can do exactly the same stuff — technology road-map, feature set — with less resource leveraging, and then move faster and charge less for it. That's a paradigm shift.

LM: What do you think of your competitors who are responding by embracing Linux and putting Linux APIs on top of their own operating systems?

BS: Well, I'm just sitting in my little office at work going, "Boy, if I was my competitor, what would I do now?" It's a logical thing to do. But that doesn't help them really. It doesn't enrich their roadmap. They are just following at that point. The more they wait the harder it gets, because the more brand value we create the more customer relationships we establish.

LM: What exactly are you making your money from right now?

BS: We sell tools.


LM: Is that where you make the bulk of your revenue?

BS: At this stage, we make the bulk of our money in services. But all of them have to do with product sales, so we still have many design wins. We're early in our company. We have a lot of design starts. Until those products actually ship and we start seeing residuals, we'll fill out our revenue stream with "hand holding." So we have three ways to make money: real product that we charge for — per unit — sometimes for our value add, tools, and professional services.

We're not a job shop. We do professional services as it relates to a product sale. If somebody comes to us and says, "The only thing I want is for you to help me do this with that Linux device; I'll never see you again," we'll say, "Go do it yourself. Good luck. You'll be back."

Robert McMillan is a contributing editor with Linux Magazine. He can be reached at bob@linux-mag.com.

May we suggest a more constructive response
to your next software consulting bill.




Visit HotDispatch. Day and night, innovative software developers are standing by to compete for your business. Save big on everything from fixes to custom projects. In a market with developer shortages and sky-high consulting costs, we're the only sane response. Ready to lower your costs dramatically? Open your free account today.

HotDispatch[®]
www.hotdispatch.com

© 2000 HotDispatch, Inc. All rights reserved. HotDispatch and hotdispatch.com are registered trademarks of HotDispatch, Inc.

THE TIES THAT



You've heard the old saying, "If you can't beat 'em, join 'em." Well, if Microsoft were to adopt a corporate motto, it might be something more like, "If you can't beat 'em, embrace and extend 'em." Seriously, Microsoft knows when it's fighting an uphill battle on a technical front. When it's clear that a Microsoft-fostered initiative will not take over the world, they eventually adopt the dominant standard. Sort of...

With Windows 2000, Microsoft finally did away with WINS (Windows Internet Naming Service, Microsoft's proprietary NETBIOS naming service) and began using the Domain Name System (DNS), the Internet's standard naming service. This seems like good news for those who run (or aspire to run) pure, simple, IP-based networks. However, Windows 2000 was written to work best with (guess what?) the Microsoft DNS Server.

Historically though, BIND, the Berkeley Internet Name Domain software, has been the most popular implementation of DNS. And while it's certainly possible to support a Windows 2000 environment using a traditional BIND name server, it requires some special configuration. In this article, we'll explore a number of configuration options and gain a thorough understanding of how to accommodate Windows 2000 with BIND.

First though, let's examine how Windows 2000 uses DNS.

Before we proceed, one quick note — for the purposes of this article, we assume that you are familiar with the inner workings of traditional BIND name servers. If you're not, please check out the article *Using DNS* in the September 2000 issue of *Linux Magazine*. You can find it on the Web at http://www.linux-mag.com/2000-09/dns_01.html.

How Windows 2000 Uses DNS

Back in the bad, old days of WINS, Windows clients "registered" their current NETBIOS name-to-address mappings with WINS servers when they were assigned IP addresses via DHCP (Dynamic Host Configuration Protocol). NETBIOS names were allocated out of a flat name space, and all WINS servers on an organization's network were supposed to have complete, synchronized copies of this name space. If a Windows client registered with one WINS server, that WINS server would tell all of its "replication partner" WINS servers about that client's name-to-address mapping. These partners would, in turn, tell their partners, and so on. This made WINS servers very chatty. This relatively unsophisticated synchronization mechanism could also leave inconsistencies among WINS servers. In addition, the use of a flat name space for NETBIOS names could result

Using BIND Name Servers with Windows 2000

in naming conflicts between WINS servers.

Now DNS provides a feature, described in RFC 2136 (<http://www.ietf.org/rfc/rfc2136.txt>), called dynamic update. Dynamic update allows an authorized updater to modify the contents of a zone on the zone's primary master name server. The updater can add or delete records from the zone and can make those operations contingent upon certain conditions, like the existence or non-existence of a particular domain name or record.

Why didn't Microsoft simply use DNS's dynamic update mechanism to handle registration, obviating the need for a separate, largely redundant naming service? Well, WINS was developed well before the advent of dynamic update.

Windows 2000 Clients

Nowadays, both Windows 2000 clients and servers use DNS. A Windows

2000 client with a static IP address registers its name-to-address and address-to-name mappings by sending dynamic updates. One update adds the client's address (A) record and another adds the client's pointer (PTR) record. For example, a client called `host.a.example` with the address `10.0.0.1` would send updates to add the two records, as illustrated in *Figure One*.

The first of these records maps `host.a.example`'s domain name to the address `10.0.0.1`. The second maps the address `10.0.0.1` to the domain name `host.a.example`. Both records are necessary to ensure correct forward- and reverse-mapping.

A Windows 2000 client with a dynamic IP address (running DHCP) also registers its name-to-address mapping, but by default its DHCP server registers the client's address-to-name mapping. In other words, the DHCP

client sends the update to add the A record, while its DHCP server sends the update to add the appropriate PTR record.

An administrator can also configure the DHCP server to handle sending all dynamic updates. This is meant for use with DHCP clients that lack the ability to update their own name-to-address mappings but works with Windows 2000 DHCP clients also. The server sets a special DHCP option that tells the client, "Just relax — I'll do all the work." Then the DHCP server sends the updates to add the A and PTR records.

Windows 2000 clients set dynamic update prerequisites when adding A records. Prerequisites allow an updater to specify conditions that must hold before the update is applied. For instance, the updater could specify that the A record for `host.a.example` that was shown earlier

Figure One

<code>host.a.example.</code>	IN	A	<code>10.0.0.1</code>
<code>1.0.0.10.in-addr.arpa.</code>	IN	PTR	<code>host.a.example.</code>

Figure Two

_http._tcp.movie.edu.	IN	SRV	0 0 80	www.movie.edu.
_http._tcp.movie.edu.	IN	SRV	1 1 80	www1.movie.edu.
_http._tcp.movie.edu.	IN	SRV	1 1 8080	www2.movie.edu.

should be added only if the domain name `host.a.example` doesn't already have an `A` record associated with it, or only if the domain name `host.a.example` doesn't already exist. In fact, this is very similar to the prerequisites actually specified by Windows 2000 clients, which are:

1. The domain name of the client must not already own an address record.
2. The domain name of the client must not be an alias.

What happens when a Windows 2000 client's update is denied because a prerequisite isn't satisfied? Well, if there's a conflicting `A` record, the client deletes the offending record and tries again. Why? Well, the `A` record could be an old record that the client added, and didn't later delete, when it had a different name-to-address mapping.

Windows 2000 Domain Controllers

Windows 2000 Domain Controllers use DNS too, but for a different reason. Domain Controllers send dynam-

ic updates to add `SRV` records that advertise where (on which host and port) various services are available. This way, when a Windows 2000 client needs to reach the Kerberos server for its domain, it can simply look up the right `SRV` record.

`SRV` records, in case you're not familiar with them, have a fairly complicated syntax, especially when compared to the more common `A` and `PTR` records. First, the owner name in an `SRV` record isn't just any old domain name; it has a special format:

```
_service._protocol.domain name
```

`service` is the name of a network service, as might appear in the `/etc/services` file, like `telnet`. `protocol` is the name of a transport protocol, usually `udp` or `tcp`. Both the service and the protocol identifiers follow an underscore, which is used to avoid collisions between domain names that own `SRV` records and domain names that represent hosts.

All `SRV` records are attached to domain names in this format. They then apply to clients of the named service running over the specified transport protocol. So, the hypothetical domain name `_http._tcp.verisign.com` applies to HTTP clients (e.g., Web browsers) running over TCP and try-

ing to access the destination domain name `verisign.com`.

Each `SRV` record has four record-specific fields: *priority*,

weight, *port*, and *target*. *Priority*, *weight*, and *port* are all unsigned, 16-bit integers, while *target* is a domain name. *Figure Two* contains a few sample `SRV` records.

Let's take *Figure Two* apart, field-by-field:

> Priority functions similarly to the preference field in an `MX` (Mail Exchange) record; if a client finds that multiple `SRV` records apply to a destination that it's accessing, it must use the `SRV` records in the order specified by the records' priority fields, lowest priority value first. So a Web browser that is retrieving these `SRV` records would first try to contact `www.movie.edu` before trying `www1.movie.edu` and `www2.movie.edu`.

> The weight field is used for load distribution; if a client finds multiple `SRV` records with the same priority, it must use the corresponding targets in proportion to their weights. So a Web browser that isn't able to contact `www.movie.edu` would then have to choose randomly between `www1.movie.edu` and `www2.movie.edu`. The client would choose `www1.movie.edu` first 50 percent of the time and `www2.movie.edu` first the other 50 percent of the time.

Figure Three

\$ORIGIN fx.movie.edu.				
@	600	A	192.253.254.14	
_kerberos._tcp.dc._msdcs	600	SRV	0 100 88	matrix.fx.movie.edu.
_ldap._tcp.dc._msdcs	600	SRV	0 100 389	matrix.fx.movie.edu.
_ldap._tcp.e437709a-1862-11d3-8eda-00400536c213.domains._msdcs	600	SRV	0 100 389	matrix.fx.movie.edu.
e4377099-1862-11d3-8eda-00400536c213._msdcs	600	CNAME		matrix.fx.movie.edu.
gc._msdcs	600	A	206.168.119.178	
_ldap._tcp.gc._msdcs	600	SRV	0 100 3268	matrix.fx.movie.edu.
_ldap._tcp.pdc._msdcs	600	SRV	0 100 389	matrix.fx.movie.edu.
_gc._tcp	600	SRV	0 100 3268	matrix.fx.movie.edu.
_kerberos._tcp	600	SRV	0 100 88	matrix.fx.movie.edu.
_kpasswd._tcp	600	SRV	0 100 464	matrix.fx.movie.edu.
_ldap._tcp	600	SRV	0 100 389	matrix.fx.movie.edu.
_kerberos._udp	600	SRV	0 100 88	matrix.fx.movie.edu.
_kpasswd._udp	600	SRV	0 100 464	matrix.fx.movie.edu.

> The port field simply tells the client which port the service is running on. HTTP, of course, normally runs on port 80. `www2.movie.edu`, however, runs its Web server on the non-standard port 8080, so an SRV-aware Web browser would try that port instead.

> Finally, the target field simply tells the client the domain name of a host running the service. The priority, weight, and port information in an SRV record only apply to the target specified in that record.

Windows 2000 Domain Controllers add a sheaf of SRV records when you set them up. *Figure Three* contains a list of the records our Domain Controller added to our zone shortly after we set it up.

Most of the records in *Figure Three* make sense if you know that a Domain Controller is basically a combo Kerberos/LDAP server; `_kerberos`, `_ldap`, and `_kpasswd` are all services that are related to these functions. `_gc` is Active Directory's Global Catalog service, which is used to find out which Active Directory server has information about a particular object.

You should also notice that two of the records aren't SRV records at all — they're A records. One is attached to the domain name of the zone ("`@`" is shorthand for "the current origin," which is `fx.movie.edu`, the domain name of our zone). Another record is attached to the domain name `gc._msdcs.fx.movie.edu`. However, all of the records — not just the SRV records — point to either the domain name or the IP address of our Domain Controller, `matrix.fx.movie.edu`.

BIND Solutions

There are several ways to configure BIND to satisfy all of Windows 2000's requirements. Not one is perfect. We'll examine these solutions and point out the advantages and disadvantages of each.

Letting Each Windows 2000 Client Send Updates

By default, each Windows 2000 client will try to update its own name-to-address mapping. If you want to permit these updates, simply allow dynamic updates to your forward-mapping zone (the zone that contains your hosts' A records) from any IP address that a Windows 2000 client might have.

There are several ways to configure BIND to satisfy all of Windows 2000's requirements. Not one is perfect; you should weigh the advantages and disadvantages of each.

First, define an access control list in the name server's `named.conf` file. This ACL will correspond to the IP addresses of all of our DHCP clients and will let us refer to them by the name "dhcp-clients" (rather than having to specify their addresses). You can do this as follows:

```
acl "dhcp-clients" {
    192.253.254/24;
};
```

Now that we've defined the ACL, we can refer to it elsewhere in the `named.conf` file. To restrict dynamic update to a zone, we use the name of the ACL in an `allow-update` substatement:

```
zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
    allow-update { "dhcp-clients"; };
};
```

Does allowing dynamic updates from any DHCP client sound scary? It is.

A user on any computer with one of those IP addresses could make just about any change to your forward-mapping zone. All he would have to do is load the appropriate software and he could make `www.fx.movie.edu` an alias for `www.whitehouse.com`. Or, he could just delete all of `fx.movie`.

`edu`. Consequently, we don't recommend this.

Letting the DHCP Server Send All Updates

A somewhat saner approach is to configure the DHCP server to handle updating both a client's name-to-address and its address-to-name mappings. You can configure this through the DHCP Manager.

Since you now know that all dynamic updates will come from your DHCP server's IP address, you can easily restrict dynamic updates to just that IP address, as illustrated in *Figure Four*.

Note that we also let the DHCP server update our reverse-mapping zone, `254.253.192.in-addr.arpa`.

What's wrong with this? Well, one problem is that the DHCP server simply believes a client when it specifies its fully qualified domain name. In other words, if you configure your Windows 2000 client with the domain name `www.fx.movie.edu`, the DHCP server will delete any conflicting address record at that domain name and add a record pointing `www.fx.movie.edu` to your client's IP address. This is no good.

Figure Four

```
acl "dhcp-server" {
    192.253.254.2;
};

zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
    allow-update { "dhcp-server"; };
};

zone "254.253.192.in-addr.arpa" {
    type master;
    file "db.192.253.254";
    allow-update { "dhcp-server"; };
};
```


Figure Five

```
acl "dhcp-clients" {
    192.253.254/24;
};

acl "dhcp-server" {
    192.253.254.2;
};

zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
};

zone "dhcp.fx.movie.edu" {
    type master;
    file "db.dhcp.fx.movie.edu";
    allow-update { "dhcp-clients"; };
};

zone "254.253.192.in-addr.arpa" {
    type master;
    file "db.192.253.254";
    allow-update { "dhcp-server"; };
};
```

Isolate Windows 2000 Clients

For maximum protection you can isolate your clients in their own zone. By default, Windows 2000 assumes that the domain name of your clients' forward-mapping zone is the same as the name of the Windows 2000 domain that they belong to. However, Windows 2000 will also allow you to specify the domain name of the zone independently.

The easiest way to do this is by right-clicking on *My Computer* and choosing *Properties*. Next, choose the *Network Identification* tab and click the *Properties* button. On the *Identification Changes* window, click the *More...* button. Now fill in the domain name of the new forward-mapping zone in the field labeled *Primary DNS suffix of this computer*. This doesn't have to be the same as the name of the Windows 2000 domain.

Let's say we create a new subzone of our main forward-mapping zone just for Windows 2000 clients. We'll call it `dhcp.fx.movie.edu` (though we could certainly configure statically addressed clients to register in it too). We would configure our BIND name servers to accept dynamic updates to

that zone from any Windows 2000 client, as illustrated in *Figure Five*.

Note that we still need to let our DHCP server update our reverse-mapping zone. If we have clients with static addresses, they will update their own PTR records. We should include them in the access control list for the reverse-mapping zone.

We've effectively created a little sandbox (or asylum, if you prefer) for our Windows 2000 clients. They can trash it if they want — deleting each other's address records, for example — but they won't interfere with records in our main forward-mapping zone, `fx.movie.edu`.

Letting Only the Domain Controller Update the Zone

Now we've got to take care of the Domain Controller. The DC, you'll remember, will try to add all sorts of SRV records (and a couple of A records) to our forward-mapping zone. We could just let it update the zone. This is less problematic than letting some random Windows 2000 client update the zone, since we have complete control of the Domain Controller. If we choose to go that route, we simply allow updates from the Domain Controller's IP address to our main forward-mapping zone, as illustrated in *Figure Six*.

We're fairly safe as long as we don't allow anyone to send arbitrary dynamic updates from the Domain Controller to our name server.

Creating Special SRV Subzones

What if you're skeptical of your Domain Controller's security? Or what if you don't trust the dynamic update code in Windows 2000 not to trash your main forward-mapping zone? Well, you can certainly take advantage of the funny format of the owner names in those SRV records and create subzones for all of those records to live in.

Remember that, generally, SRV records have the form:

```
_service._protocol.domain name
```

The protocol for most services is either `udp` or `tcp`. Microsoft also uses the pseudo-protocols `msdcs` and `sites`. So we'll just create subzones called `_udp.fx.movie.edu`, `_tcp.fx.movie.edu`, `_msdcs.fx.movie.edu` and `_sites.fx.movie.edu`. However, we'll have to take one special precaution — because the A record that the Domain Controller wants to add to the `_msdcs.fx.movie.edu` zone would other-

Figure Six

```
acl "dhcp-clients" {
    192.253.254/24;
};

acl "dhcp-server" {
    192.253.254.2;
};


acl "domain-controller" {
    192.253.254.3;
};

zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
    allow-update { "domain-controller"; };
};

zone "dhcp.fx.movie.edu" {
    type master;
    file "db.dhcp.fx.movie.edu";
    allow-update { "dhcp-clients"; };
};

zone "254.253.192.in-addr.arpa" {
    type master;
    file "db.192.253.254";
    allow-update { "dhcp-server"; };
};
```


Some Companies Don't Just Toy With Open Source.



If you've used Linux for a while, you're used to new distributions trumpeting about Open Source. They toy with the idea to be part of the herd.

Progeny Debian isn't like that. It's a contribution to the Debian project, designed to bring the strengths of Debian to a wider audience. Built by Debian's founder and a team of Debian developers, it's easier to install. It's easier to configure, and easier to find support for. But, underneath, it's pure Debian — Open Source and committed to staying that way.



Progeny Debian. Built on Debian. **Built for Debian.**

www.progeny.com

wise be illegal, we've got to turn off name checking for that zone. (Domain names that own A records are potentially hostnames. Therefore, they must follow Internet hostname syntax, which doesn't allow the use of underscores.) This would look like *Figure Seven*.

main names and records in a zone a particular TSIG key has authority to update. This would make our solution much simpler.

To whet your appetite for GSS-TSIG sup-

We will have to be satisfied with the solutions we have, inelegant as they are. At least we can use BIND to service a Windows 2000 environment.

This is a nice, safe configuration, because we can deny all dynamic updates to our main forward-mapping zone.

GSS-TSIG

How does Microsoft avoid this mess? They use "Secure Dynamic Updates," which carry with them the identity of the updater. The Microsoft DNS Server can then make sure, for example, that only the DHCP client that added a given A record can later delete it and add another A record to that domain name.

Microsoft's "Secure Dynamic Updates" are really just standard dynamic updates with transaction signatures (TSIGs), as specified in RFC 2845 (<http://www.ietf.org/rfc/rfc2845.txt>). Unfortunately, the particular dialect of TSIG that Windows 2000 clients and name servers speak, GSS-TSIG, is not presently spoken by any other name servers.

By the way, the "GSS" in "GSS-TSIG" is the same "GSS" in "GSSAPI," the Generic Security Service API. GSSAPI is an IETF standard API that provides a uniform, high-level API to security services like authentication and confidentiality.

That may change with a future version of BIND 9. BIND 9 is scheduled to support GSS-TSIG in a future release. Once it does, you'll be able to use BIND 9's new update-policy mechanism to determine which do-

port in BIND, *Figure Eight* illustrates a configuration that would ensure that only the Domain Controller could update the SRV records in the zone and that DHCP clients could only update those address records that are attached to their own domain names.

For now though, we will have to be satisfied with the solutions that we have, inelegant as they may be. What's important, however, is that you can safely use BIND to service a Windows 2000 environment.

Cricket Liu is the Director of DNS Product Development for VeriSign Global Registry Services and is author of DNS and BIND. Cricket can be reached at cricket@verisign.com.

Figure Eight

```
zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
    update-policy {
        grant *.fx.movie.edu. self *.fx.movie.edu. A;
        grant matrix.fx.movie.edu. subdomain fx.movie.edu. SRV;
    };
};
```

```
acl "dhcp-clients" {
    192.253.254/24;
};

acl "dhcp-server" {
    192.253.254.2;
};

acl "domain-controller" {
    192.253.254.3;
};

zone "fx.movie.edu" {
    type master;
    file "db.fx.movie.edu";
};

zone "_udp.fx.movie.edu" {
    type master;
    file "db._udp.fx.movie.edu";
    allow-update { "domain-controller"; };
};

zone "_tcp.fx.movie.edu" {
    type master;
    file "db._tcp.fx.movie.edu";
    allow-update { "domain-controller"; };
};

zone "_msdcs.fx.movie.edu" {
    type master;
    file "db._msdcs.fx.movie.edu";
    allow-update { "domain-controller"; };
    check-names warn;
};

zone "_sites.fx.movie.edu" {
    type master;
    file "db_sites.fx.movie.edu";
    allow-update { "domain-controller"; };
};

zone "dhcp.fx.movie.edu" {
    type master;
    file "db.dhcp.fx.movie.edu";
    allow-update { "dhcp-clients"; };
};

zone "254.253.192.in-addr.arpa" {
    type master;
    file "db.192.253.254";
    allow-update { "dhcp-server"; };
};
```


You were born with a passion

for making things work.

We'll make sure

you never outgrow it.



You've spent your entire life thinking, dreaming and creating. At DevelopOnline, we're giving you something that'll put all those ideas of yours to work a lot faster: instant access to open hardware reference platforms in an Open Source software environment. Created for next generation Web-enabled devices, these are platforms from many of today's most advanced companies. And along with them, you'll get all the development tools you need. If you think this sounds a bit revolutionary, you're right. DevelopOnline is an industry first. See for yourself. Go to developonline.com. And watch your passion for making things work grow even stronger.



DevelopOnlinesm

You're free to think againsm

GURU GUIDANCE

In the Beginning – Part II

By Aileen Frisch

Last month, we looked at the events that transpire when you boot up a Linux system, from the kernel being loaded to the *init* process getting started. This month we turn our attention to actions that follow *init*'s startup — the actions performed by the boot scripts that actually take care of all the tasks necessary to make the system ready for users.

We will begin with the system's initial boot script (as would be specified in the *bootwait inittab* entry), called *rc.sysinit* (Red Hat) or *boot* (SuSE). The script is responsible for the fundamental and essential tasks that comprise a successful boot process; it must execute successfully before the system can accept any user logins. The tasks performed by this script include (in a typical order):

- Mounting the */proc* filesystem
- Starting the kernel daemon, *kerneld*, which dynamically loads kernel modules
- Initializing RAID devices and LVM-related entities
- Running *fsck* on any filesystems to be checked
- Loading required kernel modules
- Mounting all filesystems
- Initializing basic network devices, which would include support for remote booting
- Preparing several other fundamental devices and facilities

➤ Performing various cleanup activities; these would include removing old lock files, some temporary files, lingering configuration items that are no longer needed (e.g., */etc/nologin*), and performing other similar tasks

The SuSE version of this file contains a hook that enables a system administrator to add items to the process. The relevant commands are to be placed in the file *boot.local* and will be executed near the end of the main initial boot script.

The rc Script

Once the initial boot script has finished, the boot process generally continues by executing the main boot script, *rc*, which is called with the new run level as its argument. The script functions as the overseer for the rest of the boot process; as such, it performs few actions directly itself, but rather relies on a series of other, specialized scripts to do each job that needs to be done. An idealized and simplified version of the *rc* script (derived from one on a Red Hat system) is given in *Listing One*.

Listing One: An Idealized rc Script

```
#!/bin/bash
. /etc/rc.d/init.d/functions
[ "$1" != "" ] && runlevel="$1"
if [ -d /etc/rc.d/rc$runlevel.d ]; then
# First, run the KILL scripts.
for i in /etc/rc.d/rc$runlevel.d/K*; do
# Check if the script is there.
[ ! -f $i ] && continue

# Check if the subsystem is already up.
subsys=${i#/etc/rc.d/rc$runlevel.d/K??}
[ ! -f /var/lock/subsys/$subsys ] && \
[ ! -f /var/lock/subsys/${subsys}.init ] && continue


# Bring the subsystem down.
$i stop

done

# Now run the START scripts.
for i in /etc/rc.d/rc$runlevel.d/S*; do
# Check if the script is there and subsystem is already up.
...


# Bring the subsystem up.
$i start

done
fi
```

I N T R O D U C I N G

The Most Reliable Enterprise Network Backup Management Solution.



P R O T E C T S Y O U R E N T E R P R I S E D A T A A S S E T S .

The new professional backup and restore solution for your network from Enhanced Software Technologies.

BRU-Pro is the powerful, functionally rich backup and restore utility designed for Sysadmins by Sysadmins. BRU-Pro continues the EST tradition of setting the standard you expect from your data protection software, affording you the peace of mind that the data across your enterprise is safe.

For over 16 years Enhanced Software Technologies has continued to be the leader in technology for protecting your data's integrity. We invented the concept of Backup You Can TrustSM and continue to provide the proven reliability necessary to protect your critical data.

BRU-Pro Features

- The Easiest Enterprise Backup Solution Installation
- Automated Client Recognition & Configuration
- Automated Recognition & Configuration of Tape Devices
- Allows Authorized End Users to Restore Their Files Without Tape Admin Intervention
- Tape Library Support Without Special Drivers
- Administrator Defined Security Levels
- On-line, Off-line Data Verification
- Library Bar Code Support
- Automatically Recovers From Tape Errors on Restore
- 64-Bit Computing Ready
- Heterogeneous System Support
- True Client/Server Architecture
- Servers for All Linux Variants and most popular UNIX
- Clients for Most Operating Systems
- Unattended Lights-Out Backup
- Based on The Most Awarded Linux Backup Engine



For More Details About BRU Pro visit:

www.estinc.com

From the
Developers
of the Award
Winning



The script begins by loading some functions from an external file and then sets the variable `runlevel` to the value of its first argument. The main part of the script consists of the two `for` loops enclosed within the outer `if` statement. The latter determines if a subdirectory named `rcn.d` exists in the `/etc/rc.d` directory, where *n* is the new run level. If `rcn.d` exists, the script goes on to execute the files it finds within that run level's subdirectory.

These are the files that will perform the actual tasks required to get the system up and running: initializing any remaining devices (e.g., serial/terminal lines, printers), setting up full networking, starting all the various system daemon processes for different subsystems (for example, `syslog`, `cron`, etc.), and the like.

We have yet to say much about the locations of the various system boot scripts. Traditionally, the `inittab` file is stored in `/etc`, and current Linux systems continue to follow this convention. On most Unix systems the main boot scripts — `rc.sysinit`, `boot`, and `rc` — reside in either `/etc` or `/etc/rc.d`, and Linux systems generally use the latter location. This directory also holds a series of subdirectories named `rcn.d` for each of the defined run levels (as we have just seen).

The files within each `rcn.d` subdirectory have names beginning with either

“S” or “K” followed by a two-digit number (we will discuss the significance of this in a moment), to which a short subsystem name is appended.

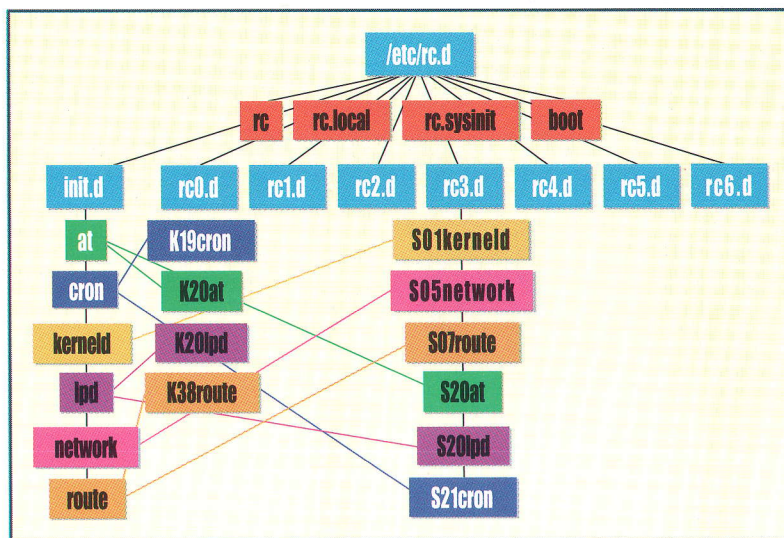


Figure One: The traditional Unix boot scripts directory structure.

However, each file is actually a link to a file in the subdirectory `/etc/rc.d/init.d` which has the same subsystem name. This directory structure is illustrated in *Figure One*.

In our example, the file named `at`, which lives in the `init.d` subdirectory, also exists (as a hard link) in the `rc0.d` subdirectory (named `K20at`) and in the `rc3.d` subdirectory (named `S20at`). Other files in the `init.d` subdirectory are similarly linked to files in these two `rcn.d` subdirectories.

In contrast, on recent SuSE Linux systems, this usual directory structure has been altered (as is also the case for some other Unix operating systems). The main boot scripts have been moved to a subdirectory of `/sbin`, in this case, `/sbin/init.d`. The `rcn.d` subdirectories also reside in this location,

and the files within them are still links to the corresponding files in an `init.d` subdirectory at their same location within the directory tree. However, the subdirectory `/sbin/init.d/init.d` does not exist but is rather a link to `/sbin/init.d` itself, and the subsystem-specific scripts actually reside there as well. Currently, `/etc/rc.d` is itself a link to the new location `/sbin/init.d`.

Keeping all of this in mind, we are now positioned to discuss each of the two loops that make up the body of the example `rc` script. Each `for` loop is constructed to run over a list of files

found in the appropriate run level-specific subdirectory and specified with a wildcard in the `for` command. For example, the first loop runs over all the files in the appropriate directory whose names began with a “K.” These files are used to shut down or stop processes and subsystems and are called “kill” files. The second loop runs over the “S” files, which are used to initiate and “start” processes and subsystem.

Starting and Killing Subsystems

At this point, it would be helpful to look at a concrete example of a directory listing from one of the subdirectories. *Example One* contains such a listing from a Red Hat Linux system.

The files within the subdirectory have the naming format that we have previously described. Note that the two-character number following the initial key letter creates an explicit sequence among the files. When the wildcard “K*” is expanded in a subdirec-

Example One: A Sample `rcn.d` Listing

```
$ ls /etc/rc.d/rc3.d
K08autofs      K20bootparamd  K30mcserver    K60atd         S30syslog      S60nfs
K09keytable    K20rusersd     K30sendmail    K60mars-nwe   S40cron       S85sound
K10named       K20rwalld      K30ypbind      K75gated      S40portmap    S91smb
K10pnserver    K20rwhod       K33ypasswd     S01kernel     S45pcmcia     S99local
K15gpm         K25innd        K35dhcpcd     S10network    S50inet
K15httpd       K25news        K35ypserv      S15nfsfs      S55routed
K15postgresql  K28amd         K40snmpd       S20random     S60lpd
```


"sangoma kicks Gluteus Maximus - I love it!"

Kenny Allman, Systems Administrator



Looking for kick-ass WAN routing? sangoma's got it.

At sangoma, our robust routing runs on whatever operating system you select. Best of all, our sophisticated WANPIPE™ software lets you monitor the health of your WAN connections, 24/7. Your network will hum while your IT department racks up the savings. Give us a call. Our solutions work....'inside the box'.

**For more information please call 1-800-388-2475
or visit our website at www.sangoma.com**

sangoma  **com**
a different 'com' company

Listing Two: SAMBA Script on Red Hat

```
#!/bin/bash

# Source function library.
. /etc/rc.d/init.d/functions

# Source networking configuration.
. /etc/sysconfig/network

# See how we were called.
case "$1" in

    start)
        echo -n "Starting SMB services: "
        daemon smbd -D
        daemon nmbd -D
        echo
        touch /var/lock/subsys/smb
        ;;

    stop)
        echo -n "Shutting down SMB services: "
        killproc smbd
        killproc nmbd
        rm -f /var/lock/subsys/smb
        echo ""
        ;;

    *)
        echo "Usage: smb {start|stop}"
        exit 1
esac
```

tory, for example, the filenames of matching files will be returned in numerical order, and the same order will be consequently followed as they are executed. When we look at the sequence of files, "K" files followed by "S" files, we can see the progress of the second phase of the boot process in detail and follow along as the subsystems are initialized one by one.

Assuming that the `for` command within the script finds some files, the loop first checks to see that the file exists. Then it extracts the subsystem name from the end of the filename (using a *bash* variable construct) and examines the `/var/lock/subsys` subdirectory for lock files related to that subsystem; this step is designed to avoid trying to stop the subsystem that is not in fact running. Finally, if these tests have all been passed, the loop runs the script, passing it the parameter `stop`.

The second loop is analogously constructed and examines and runs the "S" files in the same subdirectory, in

this case passing the parameter `start` to the script. Note that, in order to save space, we have omitted the commands for the various tests in the second loop in *Listing One*.

To summarize, on a Red Hat system, the files in the corresponding *rcn.d* subdirectory are all executed whenever a new run level is entered. Specifically, the "K" files are executed in order, followed by the "S" files in order.

Things are somewhat different on recent SuSE systems. The *rc* script that is provided with this Linux distribution first runs the "K" files in the *rcn.d* subdirectory corresponding to the previous run level and then the "S" files in the sub-

directory corresponding to the new, target run level. SuSE systems also do not use subsystem-specific lock files in `/var/lock/subsys`.

Since all of the scripts within a run-level subdirectory are links to the real script in the *init.d* subdirectory, it follows that the same script is run, regardless of whether it is executed as a "K" file or as an "S" file. *Listing Two* contains a sample script from a Red Hat Linux system; this one is concerned with the SAMBA subsystem.

This sample script again began by loading some library functions and goes on to run a networking-related configuration file stored in the `/etc/sysconfig` subdirectory. The main part of the script consists of the *case* statement, which switches on the value of the script's first argument, allowing the script to behave differently depending on whether that argument is `start` or `stop` (corresponding to being executed as an "S" file or as a "K" file, respectively).

In the first case, a message is printed to the console, and the SAMBA-related server processes are started via the `daemon` command. Finally, the lock file corresponding to the script is placed into the subsystem-specific subdirectory of `/var/lock`.

When called with the parameter `stop`, the script again sends a message to the console indicating what it is doing, terminates the Samba server processes (via the `killproc` command), and then removes the lock file.

Any parameter other than these two will cause a usage message to be displayed, after which the script will simply terminate.

All of the boot scripts on Linux systems (and on Unix systems in general) follow this general format. There are slight variations among the various scripts and across Linux distributions. *Listing Three* contains a brief excerpt from the SAMBA script (named *smb*) on a SuSE system.

Rather than loading some shell functions at the beginning, the script reads the contents of the `/etc/rc.config` con-

Listing Three: SAMBA Script on SuSE

```
. /etc/rc.config


test "$START_SMB" = "yes" || exit 0

return=$rc_done
...


case "$1" in

    start)
        echo -n "Starting SMB services:"
        startproc /usr/sbin/nmbd -D || return=$rc_failed
        startproc /usr/sbin/smbd -D || return=$rc_failed
        echo -e "$return"
        ;;
```


CONSIDER IT CORPORATE LIFE INSURANCE.



BigStorage delivers peace of mind by insuring the availability of your information. We provide nationwide support for our full lineup of storage hardware including the enterprise class tRAID with multiple host failover support, the cost-effective iRAID, and a complete NAS solution. Best of all, we do it for a fraction of the cost of the other guys. Our technical support professionals have detailed knowledge of the ReiserFS and ext2 filesystems, solid kernel configurations, and Linux friendly backup packages to meet all your mission-critical storage needs. With solutions starting under \$30,000 per terabyte coupled with our world-class 24/7 LiveResponse™ technical support, BigStorage guarantees you maximum uptime at open-systems prices.



BIGSTORAGE™

STORAGE SOLUTIONS FOR OPEN SYSTEMSSM

800-864-3789 • www.bigstorage.com

19 Heron Street San Francisco, CA 94103 Fax 415-252-1521

BIGSTORAGE AND BIGSTORAGE.COM ARE TRADEMARKS OF TORQUE SYSTEMS, INC. LINUX IS A TRADEMARK OF LINUS TORVALDS.

done
done
done
done
done
done
done
done
done
done

2

done

done
done
done

figuration file. The functions of this file are to specify which subsystems should and should not be enabled on the system and to set various characteristics for some of them. The second line of *Listing Three* illustrates this concept. It tests the value of the variable `START_SMB` and exits if that variable is set to anything other than `yes`.

Customizing the Boot Process

on is started, your system will probably die and not be available unless it is restarted manually. For this reason, it is best not to alter the ordering of these files until you are very familiar with the entire boot process and each individual subsystem.

Keeping this in mind, there are nevertheless several safe ways of customizing the boot process for the needs of a particular computer system or organization. First of all, be sure to take advantage of whatever customization facilities are provided with your Linux distribution. For example, on SuSE systems you can modify the `/etc/rc.config` file settings to enable or disable various subsystems. In addition, many of the standard boot scripts have configuration files designed to customize their functioning.

For minor additions to the boot pro-

cess you can use the built-in customization facility within the main boot script itself. Many versions of the *rc* script execute another script, conventionally named *rc.local* and stored in the same location as *rc* itself, just before exiting. Additional commands that you require can be added to this file. If your version of the *rc* script does not have such a hook, you can add one by placing a command like the following at the end of the script:


```
. /etc/rc.d/rc.local
```

You will then need to create *rc.local*. You'll also need to similarly modify new versions of the *rc* script that are installed with any system upgrades.

For more major changes, you may need to activate/deactivate, re-order, or add to the sequence of boot scripts in the run level subdirectories. This can be done manually by adding or removing properly named hard links to files within *init.d* from the appropriate subdirectory (and possibly by adding new scripts to *init.d*).

Alternatively, there are GUI tools which make this process more convenient. The best known of these is *sysvinit*; a variant created for the KDE graphical environment, *ksysv*, is illustrated in *Figure Three* (this tool is also accessible from the KDE menu).

The utility lists the contents of the *init.d* subdirectory in the leftmost pane, and the contents of the various *rcn.d* subdirectories are shown in the series of panes moving across the win-



dow to the right; "S" files and "K" files appear in separate lists. Items can often be reordered within a subdirectory by simply dragging them to the new position. Similarly, you can add a new script to any run level subdirectory by dragging the desired item from the *init.d* list to the desired location within the appropriate window. This process is illustrated in *Figure Four*, where we add (activate) a script to run-level 3 that starts the *mysql* facility.

When you drop a file into a run-level subdirectory's list, it is automati-

cally assigned the appropriate filename; the correct initial letter is chosen, and a number falling in sequence with those above and below the specified location is computed. Sometimes, however, you will need to fine-tune the utility's default numbering. This is accomplished by right-clicking on the icon whose number you wish to modify and changing it manually in the resulting dialog that displays all the items properties (illustrated in *Figure Five*).

Removing a script from one of the run-level subdirectories is equally easy.

You simply click and drag the item from the list to the utility's trashcan, located in the bottom left portion of the window.

Accept No Substitute

That's it for our Linux boot sequence overview. The addition of the GUI tools we looked at should make your life easier if you need

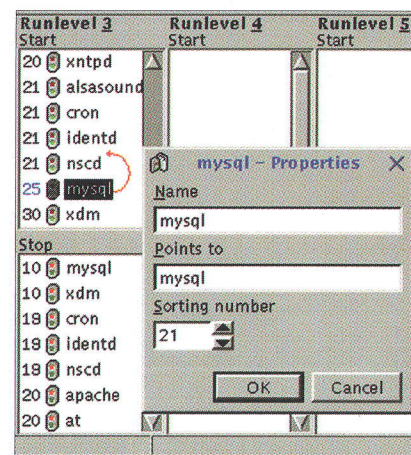


Figure Five: You can also fine tune the boot script sequencing.

to dig in and mess around with your system's boot sequence. Still, even with the addition of easy-to-use graphical interfaces, there is just no substitute for understanding what's going on under the hood.

Eleen Frisch is the author of Essential System Administration. She can be reached at aelefrisch@lorentzian.com.

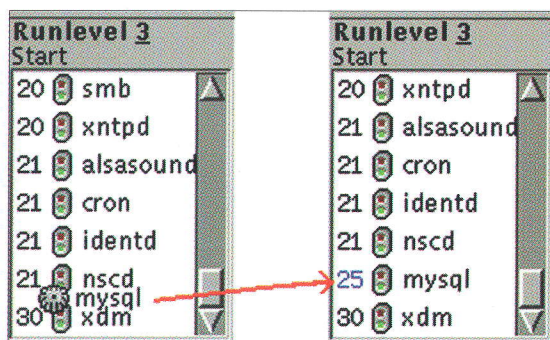


Figure Four: Adding a boot script to run level 2.

Worry Free, Private Label Web Hosting and Email?

Hostopia is the World's leading wholesaler of privately branded, front and back-end web hosting and email automation technology. Hostopia's browser based WebhostOS technology provides full automation of repetitive customer and administrative tasks including billing. Let Hostopia provide you with a user friendly, robust and profitable hosting solution **within 5 days!**

Standard features include:

SSL, ASP, MSAccess, Cold fusion, PHP, MySQL, CGI, Perl, FrontPage, CGI-BIN, Java Applets, Real Audio & Video, SSI and Web based Windows® style control panel for your customers.

- Seamlessly supports both NT and
- Totally redundant, scalable and secure
- Complete front and back-end automation

Best Hosts For Resellers™



"Hostopia.com has emerged as a pack leader in the wholesale Web Hosting and Email arena."
> Web Hosting Magazine

Your Hosting Utopia



Believe it!

Hostopia will:

- Decrease technical support costs
- Provide exceptional browser-based, website management tools
- Eliminate or reduce your research and development costs
- Automate all routine technical tasks
- Lower your operating costs
- Provide a "total solution" to increase bandwidth sales
- Reduce customer churn

Stop losing your customers to Web Hosting Specialists!
www.hostopia.com • sales@hostopia.com • 1-800-322-9438



COMPILE TIME

Threads and Mutexes – Part II

By Benjamin Chelf

Welcome to the second part of our look at programming with threads. In last month's column we talked about the functions that allow you to create and wait on threads. This month we're going to dive deeper into the problems that often arise when using threads to write concurrent programs. Before we begin that however, we'll return to the ticket agent example we looked at last month and discuss the solution to the problem of over-selling of tickets.

One quick note before we proceed: If you're writing applications that use the *pthread* libraries for Linux, you must include the option `-lpthread` as a flag to the compiler on the command line. This will tell the compiler and linker to look for the necessary functions in the *pthread* library and link them into your application.

The Ticket Agency

As we saw last month, using threads allows you to have multiple pieces of code that share and operate on a common set of data. However, threads introduce additional complexity and create code that is often difficult to debug. Let's return to our ticket agency example. Each ticket agent may be thought of as one thread. There is common data between the ticket agents — the number of tickets to be sold. Our goal is to sell all the tickets without selling too many. Each ticket agent can be modeled with the C function:

```
void* ticket_agent (void* foo)
{
    while (total_sold < NUM_TICKETS)
    {
        if (sold_ticket ())
            total_sold++;
    }
    return NULL;
}
```

Then ticket agents can be created as threads in the `main()` function, as shown in *Figure One*. This program seems simple. Ticket agents should sell tickets until they're all gone. Then, the `main()` function prints out the total number of tickets sold. However, if we actually run this program, we get some disturbing results. We need to add a few global variables:

```
#define NUM_TICKETS 10000000
#define NUM_AGENTS 4
int total_sold = 0;
```


If we run the program ten times, we will get the following:

```
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000002
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000000
mordack:~> ./a.out
10000000
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000001
mordack:~> ./a.out
10000001
mordack:~>
```

Figure One: The `ticket_agent()` Function

```
int main ()
{
    pthread_t agents[NUM_AGENTS];
    void* return_val;
    int i;
    srand (time (0));
    for (i = 0; i < NUM_AGENTS; i++)
        pthread_create (&agents[i], NULL, ticket_agent, NULL);
    for (i = 0; i < NUM_AGENTS; i++)
        pthread_join (agents[i], &return_val);
    printf ("%d\n", total_sold);
    return 0;
}
```


**Perforce Software Configuration Management
is so fast and reliable,**



you'll hardly know it's there.

You need a Software Configuration Management (SCM) tool you can rely on to do its job, so you can do yours. If it makes you wait, or gets in your way, we know you won't use it. That's why Perforce built an SCM tool that always moves as fast as you do. A careful database implementation and slight footprint ensures Perforce can turn around everyday operations in virtually zero time. Discover why over 40,000 developers rely on Perforce all day, every day. There's a free evaluation copy waiting for you. No questions. No forms. No salespeople. Simply download your copy at www.perforce.com.

PERFORCE
SOFTWARE

Fast Software Configuration Management • www.perforce.com

Notice that the program sold too many tickets eight out of the ten times! In fact, one of those times it sold two extra tickets. Automating this process and running the program 10,000 times yielded the following results:

```
mordack:~> ./a.pl
Over sold 0 exactly 3924 times.
Over sold 1 exactly 4735 times.
Over sold 2 exactly 1341 times.
Over sold 3 exactly 0 times.
```

This tells us that the program is only working correctly about 40 percent of the time. The rest of the time we are over-selling the number of tickets we have. Why?

Imagine, near the end of the program, when all but one of the tickets have been sold. One ticket agent, Bob, checks the value of the `total_sold` variable to make sure he can still sell a ticket. At that instant, the operating system's thread scheduler decides it's time for Bob to stop running and time for Sally to run for a while. Before Bob has had a chance to update the value of `total_sold`, Sally checks to make sure she can sell a ticket. Hence, both Bob and Sally can increase the value of `total_sold` when it is only one shy of the total number of tickets.

Race Conditions

Computer scientists refer to this problem as a "race condition." The problem exists because multiple threads are racing to update the values of variables based on previous values of those variables. A race condition exists any time the speed of execution of a thread, or the way threads are scheduled, can alter the outcome of the application. Race conditions are amongst the most common and most difficult problems to track down.

In our ticket agent program, the race condition exists because each ticket agent assumes that once the value of `total_sold` in the `while()` loop has been checked, that value will remain unchanged throughout the body of the loop. However, since the thread scheduler may stop the currently running thread and allow another thread to execute during that time, this is not a valid assumption.

Fortunately, the *pthread*s library provides a mechanism for avoiding this problem — the "mutex," which is an abbreviation for "mutual exclusion." A mutex allows one thread to "lock" a given piece of code so that no other thread may execute until the piece of code holding the lock is finished. In the case of our ticket agency, Bob want-

ed to make sure that once he read the value of `total_sold`, he could sell one ticket and know that he wasn't selling a ticket that no longer existed. In other words, he wanted all of the code in his `while()` loop to be executed without the possibility of some other thread executing. Bob could guarantee that the value of `total_sold` would not change if a mutex were used.

First, let's look at the *pthread*s library functions that enable you to use mutexes:

```
int pthread_mutex_lock (pthread_mutex_t* mutex);

int pthread_mutex_unlock (pthread_mutex_t* mutex);
```

These functions both return 0 on success and an error value on failure; see the manual page for more details. A thread obtains a mutex by calling the `pthread_mutex_lock()` function. When one thread has the mutex, any other thread that calls the `pthread_mutex_lock()` function will be put into a temporary state of suspended animation. Threads that are put into a state of suspended animation are referred to as "blocked."

Once the original thread calls the `pthread_mutex_unlock()` function, `pthread_mutex_lock()` will return in one of the suspended threads that called it, and that thread will gain access to the mutex. By using the `pthread_mutex_lock()` and `pthread_mutex_unlock()` functions around a given piece of code, you can guarantee that only one thread will be able to execute that code at any given time.

To create a mutex in Linux, simply declare it and assign it (in the declaration) as follows:

```
pthread_mutex_t my_mutex = PTHREAD_MUTEX_INITIALIZER;
```

Figure Two: The `ticket_agent()` Function with mutex

```
pthread_mutex_t tickets_sold_lock = PTHREAD_MUTEX_INITIALIZER;

void* ticket_agent (void* foo)
{
    int not_done = 1;
    while (not_done)
    {
        pthread_mutex_lock (&tickets_sold_lock);
        if (total_sold < NUM_TICKETS)
        {
            if (sold_ticket ())
                total_sold++;
        }
        else
            not_done = 0;
        pthread_mutex_unlock (&tickets_sold_lock);
    }
    return NULL;
}
```


There also exists a function:

```
int pthread_mutex_init
    pthread_mutex_t*
    mutex, const
    pthread_mutexattr_t*
    mutexattr;
```

you can use to initialize locks (with more information as to the type of lock and how it behaves). For our example, however, `PTHREAD_MUTEX_INITIALIZER()` is quite sufficient.

By placing a lock around the read and write of the `total_sold()` variable, each ticket agent can make sure that no other ticket agent is editing that variable. This eliminates the race condition from the program. Let's take a look at the new `ticket_agent()` function in *Figure Two*.

First of all, notice the calls to `pthread_mutex_lock()` and `pthread_mutex_unlock()`. We needed to put these calls around the code that checks the value of `total_sold` and writes to the value of `total_sold` — and we needed to put them inside the `while()` loop.

If we had placed the call to `pthread_mutex_lock()` before the `while()` loop and the call to `pthread_mutex_unlock()` after the `while()` loop body, only one thread would have been able to sell tickets. It would acquire the lock before entering the `while()` loop, therefore causing all the other threads to block before entering their `while()` loops. It would only release the lock once it had exited the loop, after all the tickets had been sold. Therefore, in the modified version, we created our own local variable (`not_done`) to track when the `while()` loop should exit, and we moved the code to lock and unlock inside the loop. This way every thread gets a chance to sell tickets.

After running this program 1,000 times, we got the following results:

Figure Three: Dining Philosophers — Part I

```
#include <pthread.h>
#include <stdio.h>

#define NUM_PHILOSOPHERS 5

pthread_mutex_t chopsticks[NUM_PHILOSOPHERS];

void think ()
{
    /* Intentionally left blank. Simulate thinking. */
}

void eat ()
{
    /* Intentionally left blank. Simulate eating. */
}

void* philosopher (void* number)
{
    int my_num = *((int*)number);
    while (1)
    {
        /* First we think. */
        think ();

        /* Grab the chopsticks to my left and to my right */
        pthread_mutex_lock (&chopsticks[my_num]);
        pthread_mutex_lock (&chopsticks[(my_num + 1) %
                                         NUM_PHILOSOPHERS]);

        /* Eat */
        printf ("Philosopher %d eating!\n", my_num);
        eat ();
        printf ("Philosopher %d done!\n", my_num);

        /* Put the chopsticks down. */
        pthread_mutex_unlock (&chopsticks[(my_num + 1) %
                                         NUM_PHILOSOPHERS]);
        pthread_mutex_unlock (&chopsticks[my_num]);
    }
    return NULL;
}

int main ()
{
    int i;
    pthread_t phils[NUM_PHILOSOPHERS];
    void* return_val;

    for (i = 0; i < NUM_PHILOSOPHERS; i++)
        pthread_mutex_init (&chopsticks[i], NULL);

    for (i = 0; i < NUM_PHILOSOPHERS; i++)
        pthread_create (&phils[i], NULL, philosopher, &i);

    for (i = 0; i < NUM_PHILOSOPHERS; i++)
        pthread_join (phils[i], &return_val);

    return 0;
}
```


Figure Four: Dining Philosophers — Part II

```
pthread_mutex_t chopstick_lock = PTHREAD_MUTEX_INITIALIZER;

void* philosopher (void* number)
{
    int my_num = *((int*)number);
    while (1)
    {
        /* First we think */
        think ();

        /* First get the lock on grabbing chopsticks.
           Then, grab the chopstick to my left and to my right */
        pthread_mutex_lock (&chopstick_lock);

        pthread_mutex_lock (&chopsticks[my_num]);
        pthread_mutex_lock (&chopsticks[(my_num + 1) %
                                         NUM_PHILOSOPHERS]);

        /* Release the chopstick grabbing lock. */
        pthread_mutex_unlock (&chopstick_lock);

        /* Eat */
        printf ("Philosopher %d eating!\n", my_num);
        eat ();
        printf ("Philosopher %d done!\n", my_num);

        /* Put the chopsticks down */
        pthread_mutex_unlock (&chopsticks[(my_num + 1) %
                                         NUM_PHILOSOPHERS]);
        pthread_mutex_unlock (&chopsticks[my_num]);
    }
    return NULL;
}
```

```
mordack:~> ./a.pl
Over sold 0 exactly 10000 times.
Over sold 1 exactly 0 times.
Over sold 2 exactly 0 times.
Over sold 3 exactly 0 times.
```

Note that we did not over-sell a ticket once in the 10,000 trial runs of the program. Although this does not guarantee that we do not have a race condition in our program, the results after the modification are much better than our original attempt.

The Dining Philosophers

Now that we've illustrated the usefulness of mutexes, let's examine a program that requires their careful usage. A famous example in concurrent programming is the Dining Philosophers program. Imagine a group of philosophers sitting around a circular table. Philosophers only do two things — think and eat. When a philosopher thinks, he has no effect on the other philosophers. When he eats however, he

can potentially create a problem for the philosophers on his left and right.

You see, there is only one plate of rice in the middle of the table and only one chopstick between each philosopher. Everyone can eat the rice, but in order to do that, a philosopher must be able to pick up his two adjacent chopsticks. Therefore, if one of the philosophers is eating, the philosophers to his right and left can't eat until he puts down his chopsticks.

We can model this problem with a program that uses threads to represent the philosophers and an array of mutexes to represent the chopsticks. When a thread holds the lock for a chopstick it is equivalent to a philosopher using that chopstick to eat. The code for the program is shown in *Figure Three* (pg. 83).

The `main()` function of this program creates all the philosopher threads and passes them a number, 0 through `NUM_PHILOSOPHERS - 1`, so that the philosopher threads know which chopsticks to grab. This program seems simple enough,

but there is a serious problem. If we run it we'll see that a lot of philosophers eat, but at some point the program will hang. It will hang at a different point each time we run it.

Deadlock

Consider the situation where every philosopher decides to eat at the same time. Because of the way we wrote this program, every philosopher will grab the chopstick to his right and then grab the chopstick to his left. However, imagine this scheduling of threads: Philosopher 0 grabs chopstick 0 and then gets swapped out, Philosopher 1 grabs chopstick 1 and then gets swapped out, and so on. Each philosopher will grab one chopstick and then wait to get the other one. However, since all the philosophers are waiting, and none are eating, they will continue to wait like that forever. This problem is called deadlock.

There is no blanket way to solve deadlock. However, there are two possible ways to eliminate the deadlock problem that exists in the Dining Philosophers program. (Why else would computer scientists like it so much?) The first solu-

tion involves using another mutex. You might have observed that the problem causing the deadlock was that a philosopher was able to grab one chopstick but then gets swapped out while another thread runs. This could be considered a race condition on the chopsticks. Each philosopher is racing to grab a second chopstick, assuming that after grabbing one chopstick, they will eventually be able to grab the other. As in the ticket agent problem, if we create a lock to make sure that only one philosopher attempts to grab chopsticks at a time, we can be sure that the philosopher will eventually get two chopsticks. The new code for the philosopher function is listed in *Figure Four*.

Now that we've guaranteed that only one philosopher can grab chopsticks at any given time, we know that the philosopher will eventually get two of them. Since he is the only one in that section, the others must be either eating or thinking.

The other solution to the problem is a bit more subtle. You might have realized that the problem with the philosophers in the original program was that they each tried to get the chopstick on the right first. If even one of the philosophers were to try to get the chopstick on the left first, the circularity that was causing the deadlock would

be broken. For this solution, we simply make half of the philosophers grab the chopstick on their right first and the other half grab the chopstick on their left first. The code for this is shown in *Figure Five*.

No Easy Answers

As we said last month, writing bug-free programs with threads is difficult. Much care must be taken to insure correct behavior. There really are no easy fixes. The use of mutexes allows you to avoid race conditions, but unfortunately, as you can see in the case of the dining philosophers, mutexes introduce the possibility of creating a deadlock in your program. Unlike race conditions, which usually cause incorrect results, deadlocks cause the program to hang entirely, which makes them absolutely no fun to debug. You now hopefully have a better understanding of the issues involved in programming with threads and the tools necessary for writing effective applications with them.

Benjamin Chelf is an author and engineer at CodeSourcery, LLC. He can be reached at chelf@codesourcery.com.

Figure Five: Dining Philosophers — Part III

```
void* philosopher (void* number)
{
    int my_num = *((int*)number);
    while (1)
    {
        /* First we think. */
        think ();

        /* If I'm an even philosopher, grab chopsticks one way. */
        if (my_num % 2)
        {
            pthread_mutex_lock (&chopsticks[my_num]);
            pthread_mutex_lock (&chopsticks[(my_num + 1) % NUM_PHILOSOPHERS]);
        }
        /* Otherwise, grab them the other way. */
        else
        {
            pthread_mutex_lock (&chopsticks[(my_num + 1) % NUM_PHILOSOPHERS]);
            pthread_mutex_lock (&chopsticks[my_num]);
        }

        /* Eat */
        printf ("Philosopher %d eating!\n", my_num);
        eat ();
        printf ("Philosopher %d done!\n", my_num);

        /* Put the chopsticks down */
        pthread_mutex_unlock (&chopsticks[(my_num + 1) % NUM_PHILOSOPHERS]);
        pthread_mutex_unlock (&chopsticks[my_num]);
    }
    return NULL;
}
```


PERL OF WISDOM

Headlines in the News

By Randal L. Schwartz

I find myself spending a lot of time participating in online discussion areas. Originally, all we had was Usenet. However, the concept of a “Web-based community” has finally taken hold. These communities usually provide some sort of message-based system (often with threading and separate discussion areas for topics) and frequently an HTML or Java-based “interactive chat” area.

I frequent one such Web community called “the Perl Monastery” (perlmonks.org). The community is active, posting dozens of messages every day, and is frequented by some sharp people, quick to answer questions. A recent posting piqued my interest. “Jcwren,” a user, suggested a series of “contests” to motivate people to show off while thinking of new solutions. He decided to kick off the first contest himself, awarding a “Perlmonks T-shirt” to the winner (funded out of his pocket).

The contest was to last only a week, and we’re midway through the week as I write this, so I can’t tell you the winner. It won’t be me, because Jcwren deliberately disallowed entries from the senior participants of the Monastery (called “saints”), of which I seem to be one.

The problem was to create a headline list for state-based headlines.

CNN’s news ticker delivers this as a far-too-flashy pop-up window.

I gave it a whack anyway. It was a nice challenge regarding a problem that’s becoming more and more common in Web-based solutions: the repackaging of information. I think we’re going to see more and more “middleware” on the Net (sites that act as brokers or meta-searchers), so I’m constantly researching to see what can be done to help.

The basic problem was to create a headline list for state-based headlines. CNN’s interactive news ticker delivers this information as a far-too-flashy pop-up window. However, the data file that it refreshes was easily reverse-engineered, and the URL and file format of that internal data file have apparently been stable for months.

Jcwren asked for a command-line program (not CGI) that fetches CNN’s internal data file each time that it is invoked. He expected this to be done from *cron* every 10 minutes or so. Any new headlines that were found there were to be

remembered in an unspecified “database.” This was to be done as simply as possible, as he wanted this to run easily on both Unix and Windows. New headlines were to be time-stamped on their first observation (there are no timestamps in the source data, so this was as close as we got to a freshness factor).

To keep the database from becoming a history book, each headline was to be aged out when it had not been seen in a specified amount of time (default one day). As long as CNN was still showing the headline, it would stay alive for at least this much longer.

To make it even more fun, the headlines had to be organized by state with a clickable set of links at the top of the output. All 50 states (plus DC) needed to always be present, but only states with current news were to be active links (which would scroll down within the document to that state). Everything had to be alphabetized, of course.

Further, the output was to be an HTML file (selectable, default `index.html` in the current directory) with a meta-refresh tag so that a browser window could be kept open on it.

I was curious about how long it would take me to write the program. I guessed around 90 minutes, and the first draft of the program was completed in just under that. I’ve since done about a half hour of tweaking. The program, which I will now describe, is present in *Listing One* (pg. 88).

Lines one through three start nearly every program I write, enabling warnings, turning on the normal compiler restrictions for large programs, and disabling the normal buffering of standard output.

Lines seven through 11 define the configurable constants that are used by this program that don’t make sense to override from the command line. The `$CNN_URL` is the source of our information. This program depends on the URL providing consistent data, so if it moves or changes format, you’re just out of luck. The `$CNN_CACHE` file is a local mirror of that remote URL. `$DB_MEMORY` holds our “database” in whatever format `dbmopen` selects (which will most often be Berkeley DB).

Line 13 pulls in the CGI module. No, this isn’t a CGI program. However, I am generating HTML so I’m using the HTML generation shortcuts, and it just so happens that the

Mini... but Mighty


Meet DEMI - Maximum Space Optimization in a Powerful Package

DEMI is the smallest and most affordable server ever designed. It's ideal for Internet based appliance system development like Linux clustering and web hosting. And it can maximize your rack space and equipment overhead savings by mounting 160 units in a standard 40U rack cabinet.


But it's not just tiny - it's powerful, too. Check out our product specs below or call us to configure it the way you want it.

If you're ready to build your own robust, ready-to-deploy appliance solution, isn't it time you met DEMI? Who else can deliver the lowest equipment overhead with the highest service revenue in the smallest space?

**2 Servers
in 1U**



Processor	Intel Celeron or Pentium III
Memory	Up to 512MB SDRAM
Network	Single or Dual 10/100 Base-T Ethernet
Hard Drive	Up to 20GB
Size	19"W x 1.75"H x 14"D
Price	starting at \$895 per node



FREE DEMI Demo

For a limited time, get a free DEMI unit on loan. For details, call **800.700.8560** and ask about our Server Appliance Development Program to help you develop the exact product your customers want.



**Siliconrax
Sliger**

Delivering ISP, Computer Telephony & Network Solutions

44200 Christy Street • Fremont, CA 94538
Toll Free: 1.800.700.8560 • Fax: 1.510.360.1610

Outside USA: 1.510.360.1600

Email: info@siliconrax.com

www.siliconrax-sliger.com

CNN input format is nearly identical to the format of an uploaded form data, which I quickly recognized to leverage off existing code. The CGI module, as of this writing, doesn't include the HTML 4.01 standard col, thead, and tbody generation methods, so I added them.

Line 14 pulls in the mirror routine from LWP::Simple (part of the CPAN-installable LWP suite).

Similarly, line 15 brings in GetOptions from the standard Getopt::Long module.

We then parse the command-line arguments in lines 17 through 23. Four variables are declared with initial values, and GetOptions alters those values if the right command-line arguments are present. You can see the GetOptions documentation for additional details (however, this should be readable as-is).

Beginning in line 25, we need a list of states. Note the

split that breaks the items on either the embedded comma or the ending newline of each line.

Lines 35 to 40 get the "current" information. Because we are maintaining a cache, we can use mirror, which minimizes the transfer cost. The request made to the server includes an "if modified since" header; if the information has not changed since this time, the server returns a quick "304 error" to tell you that you have it already. When new information arrives, the timestamp on the file is set to the "last modified" header (if present), so that the next request has the right "if modified since" header to repeat the process. Slick. Normal expected returns are status 200 (we've got a new file) and 304 (we already have the data). Anything else is broken, so we abort quickly.

Line 42 opens the database. This is a simple "hash on disk" database, so we use dbmopen to let it pick the type and nam-


Listing One: CNN Date-Based Headline Fetcher — Part I

```

1  #!/usr/bin/perl -w
2  use strict;
3  $|++;
4
5  ## begin config
6
7  my $CNN_URL = "http://headlinenews.cnn.com/
   QUICKNEWS/virtual/swf.headline.txt";
8  my $CNN_CACHE = "contest.cnn-cache"; # flat file
9  my $DB_MEMORY = "contest.memory"; # dbmopen
10
11 ## end config
12
13 use CGI qw(:all -no_debug col thead tbody);
14 use LWP::Simple qw(mirror);
15 use Getopt::Long;
16
17 GetOptions(
18     "refresh=i" => \ (my $REFRESH = 10),
19     # meta refresh time in minutes
20     "output=s" => \ (my $OUTPUT =
21     "index.html"),
22     # output file
23     "expire=i" => \ (my $EXPIRE = 1440),
24     # expire time in minutes
25     "clear!" => \ (my $CLEAR = 0),
26     # clear the cache
27     "<>" => sub { $Getopt::Long:
28     :error++; warn
29     "Unknown arg:
30     $_[0]\n" },
31 ) or die "see code for usage\n";
32
33 my @STATES = split /\s/, <<'end';
34 ALABAMA, ALASKA, ARIZONA, ARKANSAS,
35 CALIFORNIA, COLORADO, CONNECTICUT, D.C.
36 DELAWARE, FLORIDA, GEORGIA, HAWAII, IDAHO,
37 ILLINOIS, INDIANA, IOWA, KANSAS
38 KENTUCKY, LOUISIANA, MAINE, MARYLAND,
39 MASSACHUSETTS, MICHIGAN, MINNESOTA
40 MISSISSIPPI, MISSOURI, MONTANA, NEBRASKA,
41 NEVADA, NEW HAMPSHIRE, NEW JERSEY
42 NEW MEXICO, NEW YORK, NORTH CAROLINA, NORTH
43 DAKOTA, OHIO, OKLAHOMA, OREGON
44 PENNSYLVANIA, RHODE ISLAND, SOUTH CAROLINA,
45 SOUTH DAKOTA, TENNESSEE, TEXAS
46 UTAH, VERMONT, VIRGINIA, WASHINGTON, WEST
47 VIRGINIA, WISCONSIN, WYOMING
48 end
49 {
50     my $s = mirror($CNN_URL, $CNN_CACHE);
51     last if $s == 200; # we got new data
52     last if $s == 304; # no new data,
53     but we have to expire things
54     die "status is $s, aborting\n";
55 }
56 dbmopen(my %DB, $DB_MEMORY, 0644) or die
57 "Cannot dbmopen $DB_MEMORY: $!";
58 open STDIN, $CNN_CACHE or die "Cannot open
59 $CNN_CACHE: $!";
60 open STDOUT, ">$OUTPUT" or die "Cannot
61 create $OUTPUT: $!";
62 $CGI::Q = CGI->new(\*STDIN) or die "Cannot
63 parse $CNN_CACHE\n";
64 %DB = () if $CLEAR; # bye bye all that we know
65 ## first pass: add the new headlines
66 for (my $i = 1; my $headline =
67     param("headline$i"); $i++) {
68     my $state = param("state$i");
69     my $key = "$state\n$headline";
70     if (defined $DB{$key}) { # just
71         update modtime
72         $DB{$key} =~ s/\s\d+/" " . time/e;
73     } else { # add the entry



```


RackSaver™




AMD gives you the power!
RackSaver gives you the space!
With the introduction of our
UltraThin line of servers
RackSaver brings the Athlon
to the ISP and Beowulf cluster market
in the smallest chassis possible.
Athlons with DDR NOW!!
Dual processor in 'Q1 2001 for
unmatched performance. Go to our
website for instant configuration
and pricing of our systems.

Linux
P O W E R
Theirs OURS



18 nodes



84 NODES!!!



UltraThin 1100 Server

- ⇒ World's First AMD 1U Server
- ⇒ Over 1 GFLOP!
- ⇒ 1GHz CPUs available
- ⇒ Starting at just \$999.00



ThinServer 2100

- ⇒ Huge Storage & Redundant Power
- ⇒ Over 500GB per system in 2U
- ⇒ Hot-Swap drive bays available
- ⇒ Hot-Swap power supplies also



ThinServer 2300

- ⇒ 3 Systems in 2U of Space
- ⇒ Up to 63 Systems in One Rack
- ⇒ Over 100 GFLOPs in One Rack
- ⇒ Cheapest \$\$\$ per GFLOP

Now Introducing!



UltraThin 1200 Server

- ⇒ 2 Systems in 1U of Space!
- ⇒ Up to 84 Systems in a Single Rack
- ⇒ 2 drives and 1 cdrom per system
- ⇒ Ultimate redundant solution for ISP
- ⇒ Massive power for clusters

Custom Chassis Available!

RackSaver Inc.
7100 Convoy Court
San Diego, CA 92111
Ph#(858) 874-3800
Fax(858) 874-3838

WWW.RACKSAVER.COM

ing for us. (This database is cleared in line 48 if the right command-line parameter is present.)

Next, we set up the input and output streams (lines 43 and 44). The input is the file fetched from CNN. The output is the HTML file, except that we don't want to overwrite the real file just yet, so we'll append a tilde to the filename (my editor's backup file convention, so I have scripts to clean those up). After finishing the file, I'll rename this temporary file over the top of the real data in one fell swoop, so the browser will never see a partial content. This is an important strategy for uncooperative processes.

Line 46 processes STDIN using the CGI module's ability to parse a form. Assigning the output to `$CGI::Q` means that we get to use `param` and friends without having to use the ugly, and nearly always unnecessary, object-oriented form of invocation.

We then pass over the data three times. The first pass (beginning in line 50) looks for all "parameters" from the input data with the form `headline>I<n`, where *n* begins at 1 and increases (to about 100, judging from the data I saw while testing). The headline is stuffed into `$headline`, then the corresponding state jumps into `$state`.

We've now hopefully constructed a unique key of the state, a newline, and the headline. The corresponding value in the database is two integers separated by a space, both timestamps in Unix internal time format. The first number is when the headline first appeared (for display). The second number is the most recent time we've seen it (for aging purposes). So, if the key already exists, we update the second number to now, but if it doesn't, we create a new entry with both numbers set to now.

On the second pass (beginning in line 62), we age out old

Listing One: CNN Date-Based Headline Fetcher — Part II

```

58     $DB{$key} = time . " " . time;
59 }
60 }
61
62 ## second pass: expire the old headlines
63
64 for my $key (keys %DB) {
65     delete $DB{$key} if $DB{$key} =~
66         /\s(\d+)/ and $1 < time - $EXPIRE * 60;
67 }
68
69 ## final pass: generate the report
70
71 my @data
72     sort {
73         $a->[0] cmp $b->[0] or $a->[2] <=>
74             $b->[2] or $a->[1] cmp $b->[1]
75     } map {
76         [ (split /\n/), (split /\s+/,
77             $DB{$_})[0] ]
78     } keys %DB;
79
80 my %states_seen;
81
82 my @table_guts
83     map {
84         my ($state, $headline, $stamp) = @$_;
85         Tr(
86             td($states_seen{$state}++ ?
87                 escapeHTMLnobreak($state) :
88                 a({-name => fixname($state)},
89                     escapeHTMLnobreak($state))),
90             td(escapeHTMLnobreak
91                 (((localtime $stamp) =~
92                     /(.*)(\s/)[0])),
93             td(escapeHTML($headline)),
94             )."\n";
95     } @data;
96
97 print start_html(-title => "CNN Headline News",
98
99
100     head => meta({-http_equiv =>
101         'refresh',
102         -content =>
103             $REFRESH * 60}));
104
105 print table({-border => 1, -cellspacing =>
106     0, -cellpadding => 4},
107     col({-width => "0*", -align =>
108         'right'}), # state
109     col({-width => "0*"}), # date
110     col({-width => "*"}), # item
111     thead(Tr(th({-colspan => 3, -
112         align => 'center'},
113         join " | ",
114         map { $states_seen{$_} ?
115             a({-href =>
116                 fixname("#$_"),
117                 escapeHTMLnobreak($_) :
118                 escapeHTMLnobreak($_)
119             } @STATES)."\n")),
120         tbody(@table_guts));
121
122 print end_html;
123
124 close STDOUT;
125
126 rename "$OUTPUT~", $OUTPUT or die "Cannot
127     rename $OUTPUT to $OUTPUT: $!";
128
129 exit 0;
130
131 sub escapeHTMLnobreak {
132     local $_ = escapeHTML("@_");
133     s/ /&nbsp;/g;
134     $_;
135 }
136
137 sub fixname {
138     local $_ = shift;
139     tr/a-zA-Z\#/_/cs;
140     $_;
141 }

```


data by looking at all the entries' second numbers and deleting those that no longer qualify as fresh enough.

Starting with line 68, it's time to finally dump the data. For each of the keys (line 75), we pull out the state, headline, and first-seen timestamps as a three-element *arrayref*, which is then sorted by state, timestamp, and headline order.

Line 77 introduces `%states_seen`, which will be used to track the first appearance of each state in the sorted list and to figure out for which states to generate links at the top of the table.

Now comes the fun part — transforming the data into a table. First, we break each element of the `@data` array (line 89) into the three fields (line 81). Next, we create a table row (line 82) consisting of three cells (lines 83, 86, and 87). The first cell is either the state name (fixed so that it can't wrap) or on first appearance, the statename with an internal anchor. The second cell is an abbreviated portion of the localtime of the timestamp from when the headline first appeared. The final cell is the headline itself. Be particularly careful to properly encode this data as HTML entities if need be.

The next step is to generate the top of the HTML file (on `STDOUT`) handled in lines 91 to 93 with the right header, title, and meta-refresh information.

Now it's time to generate the table. The cellspacing and cellpadding are personal choices (line 95). The next three lines give hints to standards-compliant browsers (unlike

Netscape or IE) about the width and alignment of the three columns. Next comes the "table header," one row, one cell (spanning three columns), of all the states. If a state was seen, a link to the proper internal anchor is generated; otherwise, a simple name is used. Again, the state names are guaranteed not to wrap. Finally, the table guts are dumped inside the "table body" tag.

Lines 108 to 110 finish the HTML page. Once this is complete, we rename the temporary output name to its final destination with line 112.

The two subroutines starting in line 115 handle some of the needed transformations. `escapeHTMLbreak` calls the CGI-module-provided `escapeHTML` routine to fix all the HTML entities but also changes all remaining spaces to non-breaking spaces. `fixname` crunches a string so that it's a legal, unique, anchor name (for the expected dataset).

That's it. Stick it into a filepath (not in a location for your Web server's CGI and not necessarily in your `PATH`) and then run it frequently. You too will have the latest headlines from CNN. Hopefully, you'll see a few new gizmos and gadgets to steal for your own code. Enjoy!

Randal L. Schwartz is the chief Perl guru at Stonehenge Consulting and co-author of Learning Perl and Programming Perl. He can be reached at merlyn@stonehenge.com. Code listings for this column can be found at: <http://www.stonehenge.com/merlyn/LinuxMag/>

High Performance Colocation!

Performance


Multiple OC-3 Network at 5 National exchange points moves your data over the shortest possible routes.

Reliability

Redundant OC-3 connections, UPS battery backup, and an emergency generator keep your servers running optimally.

People

Our knowledgeable staff can answer your technical questions in a language you can understand - 24 / 7.



he.net

Your Industrial Internet!



HURRICANE ELECTRIC
INTERNET SERVICES

Call Today!
510.580.4190
or email
sales@he.net

TECH SUPPORT

Top Tech Support Questions

By Drew Streib

I would like to be able to save copies of all my inbound e-mail for backup purposes. How can I do this?

This is easily accomplished with a program named *procmail*. I won't go into the installation and setup of *procmail* itself, since most package systems (RPM, DEB) now do this quite nicely for you. If you already have an account on a Linux system (or your own system), chances are good that *procmail* is already setup and configured. If not, contact your system administrator or try installation via your distribution's package manager.

To quickly find out if *procmail* is already installed on your system, issue the following command:

```
# which procmail
```

A successful installation should display the location of *procmail* on your system (probably */usr/bin/procmail*).

If *procmail* is installed, the rest is easy. If there isn't one already, create a *.procmailrc* file in your home directory with permissions 0700.

```
# touch ~/.procmailrc
# chmod 0700 ~/.procmailrc
```

There are many options for *procmail* 'recipes' (the configuration sections that determine what to do with different kinds of e-mail). Some of these are explained in the man page for *procmailrc*.

```
# man procmailrc
```

For a simple start though, try the fol-

lowing in the *.procmailrc* file, substituting your own information where it is appropriate.

```
LOGFILE=$HOME/.procmail-log

:0 c:
/home/dtype/Mail/backup

:0 c:
* ^FROMdtype@dtype.org
/home/dtype/Mail/sent-mail

:0
* ^Sender: owner-linux-kernel
/home/dtype/Mail/linux-kernel
```

The LOGFILE entry specifies a file for *procmail* to log all activity to. This should be the first place you check if there are any problems.

There are three different entries in this *.procmailrc* file. The first entry specifies a file to save all mail for backup purposes (in this case, */home/dtype/Mail/backup*). The *c:* part tells *procmail* to save the mail and to continue with normal mail processing, including the rest of the configuration file. Without this, all of your mail would be saved to the backup file and nowhere else.

If you would like to save your outbound mail for tracking purposes, another entry for outbound mail is also included.

The last entry is an example of using *procmail* to sort inbound mail for a particular mailing list, in this case the *linux-kernel* mailing list. *Procmail* will look for the text *owner-linux-kernel* in the *Sender:* header of the mail. Note that the *c:* is missing here because we want the mail to be deliv-

ered to this mailbox and not to the Inbox as well.

Try sending yourself some mail to verify the proper operation of *procmail*. You should see at least a *procmail* log entry and a copy of the mail in the backup folder.

If you do not see this, then you will need to verify the *procmail* installation. */var/log/maillog* is a pretty good place to start.

This should get you started. You can do a lot with *procmail*, and it is worth learning if you need to manage a great deal of mail. At the very least, it is very useful for the backup purposes described above.

I would like to share data between operating systems on a dual-boot system. What is the recommended way to do this?

If there's one thing that Linux is good at, it's sharing. There are a lot of ways to share files from different filesystems. Some people recommend creating a separate FAT partition for sharing between operating systems, but a more elegant solution is to allow each OS to mount the other's filesystem.

Until fairly recently, there weren't any good utilities for mounting an *ext2* (Linux) filesystem under Windows. However, Linux has always been able to easily read and write to Windows/DOS FAT filesystems. Many Windows tools for reading and writing to *ext2* filesystems have emerged recently, but by and large these are far less trustworthy than the existing Linux tools for sharing filesystems.

In order to mount the FAT filesystem under Linux, your kernel will need to



PRESENTS

LINUX

BUSINESS EXPO



CHICAGO

COMMUNITY, CONTENT, COMMERCE

CONFERENCE: APRIL 2-4, 2001 > EXPOSITION: APRIL 3-5, 2001
McCORMICK PLACE > CHICAGO, ILLINOIS

FEATURED PROGRAM OF
COMDEX
CHICAGO 2001

//: OPEN SOURCE ANSWERS FOR BUSINESS.

Don't miss the Linux event designed to help you review, evaluate and compare this cutting-edge technology, and show you why Linux is best-suited for your enterprise needs.

CONFERENCE

If you're a developer or system administrator, we have the instructors and featured speakers to provide cutting-edge solutions for your Linux interoperability, integration and migration needs.

EXPO

A comprehensive expo floor that includes the LINUX Learning Center, LINUX Embedded Solutions Hot Spot, and partnerAXIS Channel Pavilion.

KEYNOTE

Tuesday, April 3, Noon - 1pm
Larry M. Augustin
President, CEO and Director
VA Linux Systems

TO LEARN MORE ABOUT
EXHIBITING OPPORTUNITIES,
call Robert Mahoney at 781-433-1629
or e-mail robert.mahoney@key3media.com.

//:REGISTER ONLINE TODAY AT WWW.LINUXBIZEXPO.COM

Figure One: Configuring Apache for Wake One Semantics

```
# CFLAGS='-DSINGLE_LISTEN_UNSERIALIZED_ACCEPT' ./configure
```

contain the appropriate support. Most default kernels already include this support, but if you are compiling your own kernel, be sure to include FAT and VFAT filesystem support. If you'd prefer not to compile it into your kernel, FAT and VFAT support are also available as loadable modules.

In any case, to start sharing filesystems you'll need to first create a mount point under Linux. This can be any directory in the filesystem that you want, but it's helpful to use a name that makes sense (such as `/mnt/msdos` for this example).

```
# mkdir /mnt/msdos
```

You will then need to add a line to your `/etc/fstab` for the FAT filesystem. If your kernel contains VFAT support, you'll be able to view and edit the long filenames used by modern versions of Windows.

Try adding the following line to `/etc/fstab` (note that in this example, we used the third partition `/hda2` on our hard disk — you will, of course, need to substitute the proper partition for your system):

```
/dev/hda2 /mnt/msdos vfat defaults 0 1
```

If you don't want this filesystem to be automatically mounted at boot time, you can add a `,noauto` after the `defaults` option. You should now be able to mount the Windows partition as the root user.

```
# mount /mnt/msdos
```

If you get any errors about a bad filesystem, then you have likely specified the wrong partition. At this point, you should probably verify the partition and retry.

If you get an error regarding kernel support for the `vfat` filesystem, try

changing the `/etc/fstab` entry from `vfat` to `msdos`. If this works, then your kernel doesn't have support for Windows' long filenames. If this does not work, your kernel does not have the proper filesystem support for any DOS or Windows filesystems.

Assuming that everything *did* work, however, you should now be able to access your Windows data from Linux under the `/mnt/msdos` mount point (or whatever mount point you created). If you set this up as shown above, only root will be able to write to these files. Keep in mind that the FAT filesystem does not include user permissions support.

I'm trying to service a lot of connections via Apache. I have hard configured the number of Apache children to 1000 in order to minimize process spinup time. The box performance is bad and occasionally even refuses connections. Any ideas?

You're likely running into a problem often referred to as the "thundering herd." It is well documented with Apache and Linux, and there are several ways to work around it.

In Linux, this condition usually stems from the process "wake semantics." When a new connection arrives to be serviced in Apache/Linux, all of the sleeping processes are notified. All of these processes will then all try to take control of the new connection. However, only one of them will succeed, and the others will fail and return to their sleeping state. This is referred to as wake all semantics. Linux 2.2 and older kernels operate in this manner.

When there are only a few sleeping Apache processes, this isn't a problem. Apache normally will scale the number of sleeping processes using the

MinSpareServers and MaxSpareServers configuration variables from `httpd.conf`. Setting the MaxSpareServers to an abnormally high value, however, can potentially cause performance problems. I usually set it to between five and ten percent of the MaxClients value.

If you are setting the MaxSpareServers value abnormally high, tuning it down may result in an immediate performance gain.

Some kernels don't suffer from this problem because of an ability to utilize 'wake one' semantics, which allow a single process to be awakened for each inbound connection. The BSDs are capable of this, as is the Linux 2.4 kernel.

In order to take advantage of these semantics, Apache must be compiled with a particular option. When configuring Apache for the 2.4 kernel, try issuing the command that is shown in *Figure One* before compiling. This may improve performance on kernels capable of wake one semantics.

There are a variety of tuning parameters for servicing high numbers of network connections. Apache has a built-in hard limit for the maximum number of connections allowed, requiring a recompile to set this higher than 256.

Linux will also require some `/proc` filesystem tuning in order to accommodate high-capacity servers. You will also want to look at Linux's compiled-in process, system-wide, and per-user limits.

In short, there are a lot of things to be aware of when tuning for high numbers of Apache processes/connections. The "thundering herd" condition is perhaps one of the most overlooked problems, but there are a variety of other factors to consider.

More information about tuning Apache for high numbers of connections can be found at <http://linuxperf.nl/linux.org/webserving>.

Drew Streib is a coder, admin, and writer with VA Linux Systems. He can be reached at tech@linux-mag.com.

All Internet. All Business. All Week.

The World's Largest Event

Dedicated to e-Business

and Internet Technology

Featuring:

- The industry's only five-day conference on emerging Internet applications, technologies and trends.
- Brand new forums on the Urban Internet, ASPs, the e-business workforce and more.
- Enlightening keynote presentations from five top industry leaders.
- More than 900 exhibitors and the world's largest display of Internet products and services.
- Co-located with eCRM, the first event to focus exclusively on strategies for customer relationship management.

Internet World Spring Presents:



Los Angeles Convention Center
D²ME — Digitization & Distribution of Media & Entertainment consists of two days of conference forums.



An online competition of media and entertainment. Live-action films, animation, web commercials, web series and experimental works. Cast your vote online beginning February 7 at www.internetworldmefestival.com.

Register Now!

at www.internetworld.com/registerspring

For More Information

Visit www.internetworld.com

To Request a Brochure

Call 1-800-500-1959 or e-mail iwsprogram@iw.com

To Exhibit

Call 203-559-2884 or e-mail twalsh@iw.com

The Eighth Annual Internet World Spring 2001

Conference: March 12-16
Exhibits: March 14-16

Los Angeles Convention Center
Los Angeles



INTERNET WORLD
Spring 2001

Produced by:



Official Sponsors:



INTERNETWORLD.com

Media Sponsors:



internet.com

m-CommerceWorld.com

Zonker's Product Picks

By Joe "Zonker" Brockmeier

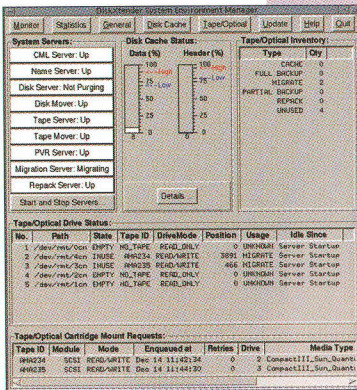


DISKXtender 2000: Infinite Storage

It's commonly believed that you can never be too rich or too thin. This isn't always true; just look at Calista Flockhart. Money and skinniness aside, one thing you truly can't have enough of is data-storage space. No matter how much space you have, it'll be filled eventually — especially in a production environment, especially at the end of a fiscal quarter.

To the rescue comes DISKXtender for Linux. DISKXtender doesn't actually buy more storage space, but rather allows for the seamless integration of additional storage as needed. DISKXtender also enables distributed data access and reduces the overall cost of data management.

DISKXtender supports optical, tape, and DVD storage media. It is sold on a tiered pricing plan, based on storage capacity, and starting at \$10,000, is a bit pricey. It's a product of OTG Software, based in Bethesda, MD. <http://www.otg.com>



Mozilla 0.7: Better Browsing

Browsing on Linux is getting good. The latest build from the Mozilla project is fast, stable, and looks pretty. Netscape went open source nearly two years ago, and it appears that the Mozilla project is finally starting to bear fruit. The latest Mozilla build seems much more stable than even the 4.7 line of Netscape browsers, and the recent Netscape 6 release is based on Mozilla 0.6.

Mozilla touts the best W3C standards support of any browser and allows the user to customize its appearance using themed skins, which are available from a number of sites.

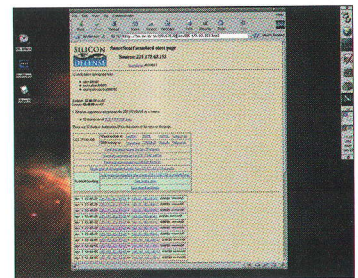
Grab the latest Mozilla build from Mozilla Project's Web site at <http://www.mozilla.org>.



Snort: Sniffing Intruders Out

Snort is a must-have program for the security conscious. The intrusion-detection program performs real-time packet logging and traffic analysis on IP networks. It's used to detect attacks and probes on IP networks, including port scans, buffer overflows, attempts to detect a host's OS, and many other intrusion schemes.

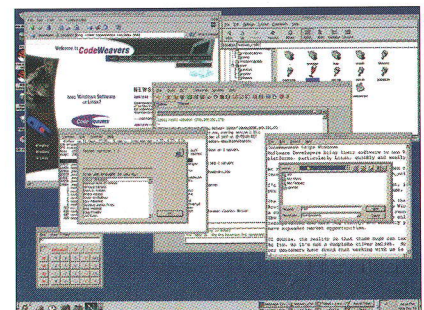
Snort runs on a variety of operating systems and platforms, including Linux, *BSD, SunOS and Solaris, AIX, and even Mac OS X Server. Snort is distributed under the GNU Public License and is written by Marty Roesch and a number of other contributors. <http://www.snort.org/>.



CodeWeaver's Wine: 2001 Is a Good Year for Wine

The Wine project has been around for quite some time. However, installation and configuration has generally been beyond the capacity of mere mortals. Wine is designed to run Windows programs under Linux, allowing migrating users to utilize their favorite applications without switching back and forth between the two. While Wine isn't quite there yet, it's pretty darn close, and now many Windows apps run well on it.

CodeWeaver's release of Wine features a number of enhancements that make it easier for non-gurus to install and use. CodeWeaver's Wine comes with a configuration wizard that's friendly and fairly easy to use. Once you've installed Wine, the Wine Launcher utilizes a friendlier interface for presenting status and diagnostic information. For more information and downloads, visit CodeWeaver's Wine site at <http://wine.codeweavers.com>.




BlackAdder: Portable Python

Python developers have a new tool to add to their arsenal. The Kompany's BlackAdder Python IDE will slither its way into their hearts. BlackAdder is a full-featured IDE for Python that includes syntax highlighting, cross-platform ODBC interfaces, and a debugger.

BlackAdder comes in a Linux and/or Windows version, so companies with mixed development environments can provide their Linux and Windows developers with the same interface. That's not all folks! BlackAdder also generates cross-platform code, so Python apps developed with BlackAdder run on Windows or Linux.

The BlackAdder IDE is available for personal use at \$79.99 and for business use starting at \$399, which includes both Linux and Windows versions. <http://www.thekompany.com/products/blackadder>.



the **KOMPANY**.com



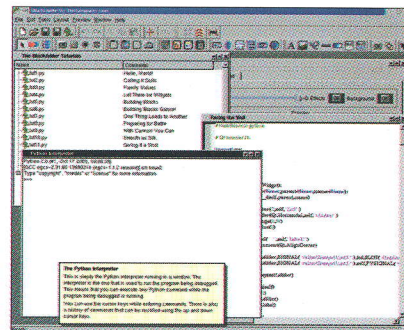
stormix
TECHNOLOGIES INC.

Storm Strike 1U Server: Servin' Up a Storm

Storm's Strike Server is more than just a sexy bit of hardware. It's an easy to administer server that spares space in the server room.

Strike comes preloaded with Stormix, Storm's version of Debian GNU/Linux. Configuring your Strike server is easy enough. You can use either the LCD control panel on the unit or configure it remotely through the Stormix Web interface. The Strike Server is a 1U rack unit that's available in a number of configurations, including single and dual CPU, IDE or SCSI disks, up to 1 GB of RAM, and up to three Ethernet NICs.


Stormix Technologies is based in Vancouver, British Columbia. You can find out more about Stormix and the Strike server at <http://www.stormix.com>.



Gimp 1.2: Graphics Goodness Galore

In case you hadn't heard, there's a new GIMP in town. Version 1.2 was released with even less fanfare from the developers (if that's possible) than the 2.4 kernel.

For the uninitiated, the GIMP is the GNU Image Manipulation Program. It's a Photoshop clone for Linux that is great for creating Web graphics and touching up photos and other graphical works. The GIMP comes with a number of plugins for image manipulation and opens nearly every image format. While the GIMP is easy to get started with, it also has many advanced features that, once mastered, will yield spectacular results. The GIMP is distributed under the GNU Public License. To grab the latest version of the GIMP, check out the GIMP homepage at <http://www.gimp.org>.




VISUAL
SlickEdit®

Visual SlickEdit: Edit in Style


MicroEdge's Visual SlickEdit is loaded with features for programmers who need a little more oomph in their editor. Visual SlickEdit features support for just about every FTP client, allowing you to manage your files nearly transparently via FTP. SlickEdit also emulates vi, Emacs, and even Visual C++ for programmers who want to make the switch but are afraid of losing their "touch." SlickEdit is full-featured, having tons of bells and whistles and allowing for hex editing, syntax highlighting, three-way file merge, and drag-and-drop editing. A new feature to SlickEdit 6.0 is DIFFzilla, an advanced differencing tool that allows for viewing and merging changes from two files, directories, or source trees.

SlickEdit is a product of MicroEdge and is available for Linux, Solaris, AIX, HP-UX, IRIX, Windows, and several other platforms. SlickEdit starts at \$295 for Linux and is available at <http://www.slickedit.com>.



Got a product pick of your own for Zonker? Drop him a line at zonker@linux-mag.com.

Advertiser Index



To advertise in *Linux Magazine*, please contact the appropriate ad sales rep listed below for a 2001 media kit containing an editorial schedule, rate card, and ad close dates.

East Coast Sales

Eric Derby • 973.786.0340 • eric@linux-mag.com

West Coast Sales

Mike Schaps • 650.966.8476 • mike@linux-mag.com

ASL Workstations . . .	http://www.aslab.com	21
Big Storage	http://www.bigstorage.com	77
Caldera	http://www.calderasystems.com	5
Chilliware	http://www.chilliware.net	43
Consensys	http://www.raidzone.com	25
Cyclades Corporation .	http://www.cyclades.com	9
DevelopOnline.com . .	http://www.developonline.com	71
Einux	http://www.einux.com	31
eLinux	http://www.elinux.com	51
Enhanced Software Technologies	http://www.estinc.com	73
Hewlett Packard Co. .	http://www.hp.com	12
Hostopia.com	http://www.hostopia.com	79
HostPro	http://www.hostpro.com	56
HotDispatch.com . . .	http://www.hotdispatch.com	63
Hurricane Electric . .	http://www.he.net	91
InternetWorld Spring .	http://www.internetworld.com/spring2001/	95
Key3Media Events . . .	http://www.linuxbizexpo.com	93
Knox Software	http://www.arkeia.com	27
Linuxcare	http://www.linuxcare.com	2
Linuxgruven	http://www.linuxgruven.com	35
Microway	http://www.microway.com	19
NeTraverse	http://www.nettraverse.com	36
Network Shell Inc.	http://www.networkshell.com	62
NPI	http://www.npix.com	29
OTG Software, Inc.	http://www.otg.com	55
Penguin Computing	http://www.penguincomputing.com	22
Perforce Software	http://www.perforce.com	81
PFU America	http://www.pfua.com	44
Progeny Linux Systems . .	http://www.progeny.com	79
RackSaver Inc.	http://www.racksaver.com	89
Rackspace.com	http://www.www.rackspace.com	45
Red Hat	http://www.redhat.com	C4
Sangoma Technologies . .	http://www.sangoma.com	75
Siliconrax-Sliger	http://www.siliconrax-sliger.com	87
Software Development Conference & Expo	http://www.sdexpo.com	99
SuSE	http://www.suse.com	7
Terra Soft Solutions	http://www.terrasoftsolutions.com	54
Turbolinux	http://www.turbolinux.com	C3
VA Linux Systems	http://www.valinux.com	C2
ValueWeb	http://www.valueweb.com	33
Verio Inc.	http://www.dedicatedsever.com	53
Veritas Software	http://www.veritas.com	15
Western Telematic, Inc. . .	http://www.wti.com	17

Linux Magazine is published monthly by InfoStrada LLC at 234 Escuela Avenue #64 Mountain View, CA 94040. The U.S. subscription rate is \$29.95 for 12 issues. In Canada and Mexico, a one-year subscription is \$39.95 US. In all other countries, the annual rate is \$49.95 US. Non-US subscriptions must be pre-paid in US funds drawn on a US bank. Application to Mail at Periodicals Postage Rates is Pending at Mountain View, CA and Liberty, MO. POSTMASTER: Send address changes to *Linux Magazine*, P.O. Box 55731, Boulder, CO 80323-5731.


Article submissions and letters should be e-mailed to editors@linux-mag.com. *Linux Magazine* reserves the right to edit all submissions and assumes no responsibility for unsolicited material. Subscription requests should be e-mailed to linuxmag@neodata.com or visit our Web site at www.linux-mag.com.

Linux® is a registered trademark of Linus Torvalds. All rights reserved. Copyright 2001 InfoStrada LLC. *Linux Magazine* is printed in the USA.



**Software Development
Conference & Expo
invites you to...**

FEED YOUR HEAD



The SD Conference and Expo is chock-full of essential ingredients for the professional developer. Only SD offers a curriculum that is wide, deep and relevant to help solve your daily development challenges. From languages like Java, C++, XML to cross platform integration to architecture, design and management - SD has it all.

For the complete menu, visit our website at www.sdexpo.com or call 1.800.441.8826.

Visit sdexpo.com
for more information.

*In an independent research study by Evans Marketing Services for the year 2000, the Software Development Conference and Expo was rated the number one conference for software development teams.

PREMIERING THIS YEAR

WEB SERVICES WORLD

A two day conference-within-a-conference devoted to the open standards of service-oriented architecture in support of the UDDI initiative.

**Don't miss this
ground-breaking event!**

**Coming April 8 - 12, 2001
San Jose, California**

Serious content for serious developers

SHUTDOWN

**DarkSide Dave
For Columbia Internet:**



Is no doubt, I am
Lord of Legion of Evil.
Am havink many wery
obedient minions, who
am trickink to leave jobs
in Nike factories.

They are workink hard,
but am always lookink for
ways to make more cash
money with talented but
foolish mortals,

So am usink Linux. Is
like bait for geeks.
They are thinkink,
"He is usink Linux, he
must be wery smart
and wery nice man.
I will work for him."

Of course, by this time,
have already chained
geek to oar.

(Needink another 42
geeks. Triremes, they are
hard to fill with oar slaves.)

LINUX

"Where is your despotism?"



How to achieve double-digit productivity gains and seven-digit cost savings:

Now you can get supercomputing power without the super-sized price tag. TurboLinux software can make your existing IT infrastructure work harder by creating high-performance clusters with the equipment you already own, whether you're running Linux, UNIX, or Windows NT/2000. Our clustering technology has helped financial powerhouse J.P. Morgan harness the idle CPU cycles on hundreds of Windows NT workstations. They're running complex derivatives models faster—and they're saving \$7 million a year. Easy to manage and completely scalable, TurboLinux solutions offer one more thing: peace of mind. Our knowledgeable trainers and consultants can help you get started and keep your enterprise running at peak efficiency.

How to learn more:

In the U.S., give us a call at (650) 228-5093. Or visit us online at www.turbolinux.com/learnmore/savings.html.




turbolinux.
Powerful Thinking.



redhat
CERTIFIED
ENGINEER

128-bit Encrypted Job Security.

Become a Red Hat Certified Engineer—the most popular, most respected Linux certification. The choice of IT pros.

You know no job is forever. Bosses leave, departments change, budgets disappear. But with the right skills you can make yourself invaluable.

Red Hat Certified Engineer™ tests your ability to install and configure Linux®, as well as file systems, security, and open source applications like Apache and Sendmail.

Prepare through intensive hands-on classroom instruction or via eLearning. We'll teach you today's most critical network services. We'll test you on live equipment.

When you succeed, the skills, job security, and the confidence that goes with it are all yours. Along with the best four letters of recommendation you'll ever earn: RHCE™.

www.redhat.com/training



redhat.

Good guys wear red hats.